2009 COV Executive Overview

The 2009 CISE/CNS Committee of Visitors (COV) reviewed the CNS program portfolio during the period 2006-2008. This report reviews findings and summarizes recommendations. Additional details can be found in our cluster subgroup reports and the CNS self-study, on which we relied heavily for responses to the statistical questions posed in part A of the COV template.

In broad summary, our COV was impressed by the scope and quality of the work underway in the CNS division, felt that programs are well managed, and was convinced that the scientific impact of the research CNS is supporting is substantial. We felt that CNS is positioned to have even greater impact in the future as the country begins to develop cyber-based technologies for critical sectors such as medical, financial and power. CNS is to be commended for its high standards, efficiency and impact.

Our most serious concern relates to the inadequacy of resources for program managers (for example, these individuals are not currently funded to attend the top-ranked conferences in the areas of their own programs, a serious deficiency in our view). We also find that CNS has been more aggressive about revamping programs than tracking the impact of programmatic changes on the research community, and as a result that it is often difficult to understand whether those changes had beneficial impacts, negative ones, or some mixture. We elaborate on these and other issues we identified and make recommendations below.

The first section of the report briefly lays out background for the time period under review. The next section discusses our major findings. A subsequent section gives recommendations. We then provide the requested statistics on COV makeup, agenda, and proposal sampling. Last, we include the detailed COV reports from our clusters. We include but did not prepare the CNS self-study.

Setting the Context

Before laying out our findings, it will be useful to review some aspects of the context in which the reporting period falls and the associated challenges faced by the CNS programs.

Context: Relatively Flat NSF budgets; Research Institutions Under Stress.

During the reporting period, NSF as a whole had relatively flat budgets. CNS (and CISE as a whole) received record numbers of proposals, and the overall CNS budget rose 23.5%. Obviously, there was variability within individual programs: some CNS programs were cut sharply (for example, NeTS rose from $34M in 2006 to $40M in 2007, but then dropped to $29M in 2008), while others were launched (for example, the $12M GENI effort).

During the reporting period, DARPA pulled out of academic research funding, with impacts on the research community that have been studied by others (notably the Computing Research Association[1] and the National Academy of Sciences[2]). It was not possible for us to isolate the specific impacts this may have had within CNS, although we do think it is obviously a contributor to the growth in numbers of proposals that CNS received during the 2004-2006 period (numbers stabilized at a fairly high level during the 2006-2008 period that we studied here).

As the reporting period ended, the severe financial downturn impacted endowments and state funding streams, and sharply depressed technology sector performance. These events are reducing state and corporate research funding streams, forcing researchers to look to NSF for a larger percentage of their overall research funding. On the positive side, if CNS receives recovery act funding, those funds could be used (in part) to ease funding pressures within the programs.

Relative to the prior COV period, proposal success rates improved during our reporting period, approaching norms for NSF as a whole. In particular, the previous COV self-study indicated that NSF achieves a 23-27% success rate for proposals, over all directorates. During 2003-2005, CNS success rates were closer to 16-21%. During our reporting period, success rates averaged 26.3%. As noted above, CNS funding also rose by 23.5% from 141M to 174M/year. The COV was pleased by this progress.

---

[1] For example, see "An Endless Frontier Postponed", Edward D. Lazowska and David A. Patterson, Science 6 May 2005: Vol. 308. no. 5723, p. 757.

2 *Assessing the Impacts of Changes in the Information Technology R&D Ecosystem: Retaining Leadership in an Increasingly Global Environment.* National Academy of Sciences CSTB, 2009. http://www.nap.edu/catalog.php?record_id=12174

Context: An Approaching Period of Dramatic, Disruptive Change

Our COV convened just as a wave of tremendously disruptive change has gained force, positioning CNS to have greater potential impact over the coming five years or so than at any time since the emergence of the web in the early 1990s. Any of a number of developments, including cloud computing, multicore processors, virtualization, trusted computing, and true mobile computing with cellphone platforms, has the individual potential to revolutionize the area. These purely technical drivers are further complicated by the explosive growth of social networks and the associated science, which are driving a sweeping change in the ways we use computing.

Also important is the emergence of new kinds of societally important application areas. Medical informatics, long a sleepy backwater, is poised to experience a tremendous shift in vitality and criticality as the government presses for change in the health care industry. Control of modernized power distribution systems and other civil infrastructures may be of equal importance as we move to improve efficiency in response to concerns about the global climate. Taken jointly, such developments will change the nature of the questions we need to explore in such areas as Cybertrust and NeTS.

Moreover, we are doubtful that industry, in and of itself, can solve these questions. For example, Internet hosting systems like Google, Yahoo! and MSN Live work by shipping data to massive data centers for indexing, linking, and collation. However, security and privacy concerns make it difficult to imagine shipping the nation's medical data to any one place. We see the CNS community as a strong candidate for creating new models for systems that address these concerns.

Considered in aggregate, such developments set the stage for CNS to have massive impact over the next 5 years, much as has been the case with other major disruptions in computing history—assuming, that is, the CNS community engages the opportunity in a sophisticated way and designs programs that rise to the challenge.

Context: A CyberSecurity Crisis

The period under discussion was also shaped by a steadily growing cybersecurity crisis. An arms race erupted between the community attacking the cyber-infrastructure and the community growing that infrastructure. Many government-sponsored studies confirm that the attackers' "black market" is booming. Whether for disseminating spam, mining for industrial secrets, government espionage, attempts at blackmail, or as a vector for malicious hackers bent on inflicting harm, this market generates substantial revenue. Simultaneously, various nation-state "actors" have become significant players in the area of cyber-espionage. In 2007-2008, even some well-managed, highly protected U.S. government installations reportedly were victims of sophisticated intrusions.

Not surprisingly, Congress has identified cybersecurity as a national priority, and CNS has become the flagship organization in the area.

Given the rolling nature of the threat, this is not a battle that can be “won” in the foreseeable future, but it is certainly possible to reduce the risks associated with entire families of attacks. For example, so-called synthetic diversity mechanisms deployed in some modern operating systems use randomization to defeat a large class of vulnerabilities such as buffer overflow attacks. This concept has yielded dramatic improvements in platform security and is now used commercially in Windows Vista and some versions of Linux, the two most widely used commercial software platforms.

DISCUSSION OF MAJOR COV FINDINGS

Broadly, our COV believes that overall CNS funding levels continue to be less than what is needed. As noted in the CRA and NAS studies, the information technology R&D community has been a powerful engine for economic growth and development over an extended period. The country confronts important challenges as we move to exploit information tools in settings such as healthcare and management of the power grid, which pose questions of privacy and robustness unlike those that arise in the most widely deployed commercial computing platforms. CNS is positioned to have a huge impact on such problems, provided that it has the funding and broad directives to do so. Much evidence supports the view that investments in these technical areas can catalyze US economic growth, stimulate high-tech job creation, and enhance national competitiveness.

More specifically:

- Part A of the COV charge concerns metrics for program quality and breadth, in multiple senses. The COV was satisfied that CNS is performing above our expectations and, in many cases, substantially exceeding our expectations.
- We identified some cases in which CNS seems to be stretching limited funds and raising proposal success rates in part by negotiating budgets and timelines. This can be appropriate if a budget is genuinely too high for the work proposed. However, it can also harm scientific work and broader impacts, and in the long run reduce scientific output per dollar invested. The tradeoff needs continuous monitoring and adjustment.
- The prior COV found evidence that high rejection rates were leading at least some researchers to submit multiple proposals to multiple programs, and to resubmit rejected proposals, strategies that tend to amplify the numbers of pending submissions. Our COV was unable to determine the extent to which this phenomenon continues.
- A previous COV recommended that CNS limit the numbers of proposals individual researchers could submit and the frequency at which new proposals would be considered. CNS must monitor the impact of such policies, to recognize and address any deleterious consequences.
- A previous COV expressed concerns about CNS panel quality and diversity. Our COV concluded that these have been resolved in a satisfactory manner.
- A previous COV worried that reviewers lack a mechanism for reporting their own confidence in reviews; CNS is implementing a satisfactory mechanism at this time.

CNS program managers brought to our attention a puzzling phenomenon: they pointed out that CNS panels appear to be overly harsh in their proposal “rankings”. A consequence is that genuinely outstanding work sometimes doesn’t receive the highest rankings. The COV sampled reviews and concluded that the CNS program managers are correct in this respect.

- We believe the issue can be traced to a community trait: conferences in the systems area tend to favor aggressive styles of paper reviews. Within our top conferences, program committee members are sparing in their use of the highest rankings. We suspect that CNS panelists, drawn from the same population that participates in these conference program committees, may simply be bringing this prior experience to the table when asked to review for CNS.
- Were the issue confined to CNS, it might not matter very much; if the community favors the term "very good" under conditions in which other research communities might have used the term "excellent", one could call it a matter of taste. However, CISE program officers have made it clear that if ratings cannot be compared across all of NSF's directorates in a consistent manner, confusion arises, and CISE is placed at a disadvantage when competing for funding with other NSF directorates.
- Accordingly, our COV accepts that the issue is more than just cosmetic. Given the goal that CISE proposal ratings be undertaken on the same scale used by other directorates, we believe that the best strategy is to focus on educating panelists and external reviewers about the NSF rating scale and how it is normally used. With this in mind, our COV recommends that CNS make NSF-wide statistics available to reviewers, to show them the typical percentages of proposals receiving each rating in other directorates (that is, outside of CISE). Program managers should also explain the issue to panelists, so that they will understand that a CISE panel is not identical to a program committee in this respect. One might hope that a better educated reviewer community would use ratings in a manner more consistent with those in the remainder of NSF.

CNS panels as a whole impressed us for quality of panelists, diversity, and quality of work.

- We note, however, that individual reviews continue to be of variable quality (ranging from a single line in some cases, to short essays in others).
- CNS panel summaries were also felt to be of variable quality although, in most clusters, more consistently good than reviews. Some panel summaries were clearly written using cut-and-paste from whichever review was felt to be most complete.
- On the other hand, the proposal feedback provided by program officers was highly impressive, showing tact, balance, and sophistication. We do not feel that it is appropriate for program officers to be put in a position of needing to compensate for work that panelists should be doing, but they must be commended for their outstanding performance in the role.

With respect to part B (impact), we were extremely impressed.

- CNS has maintained a consistent track record of funding some of the most important computer science research underway in the nation today.
- We were hard pressed to decide which nuggets to include into this report, because there were so many to chose from.
- CNS-funded research work is impacting the science base and the commercial sector in measurable, vital ways and responding to urgent national priorities in cybersecurity.
- CNS is achieving these objectives in an increasingly complex funding landscape, stepping in to fill gaps despite a serious shortage of resources – and succeeding even in the face of tough odds.

The "education and workforce" cluster of programs should be applauded for its impact with respect to promoting NSF's mission with respect to "Learning", developing the science and engineering workforce, and promoting scientific literacy among the broadest possible spectrum of society. Howevery,

- Program officers in this area are having a particularly difficult problem with reviewers and review quality, putting them under load and stress.
- The REU program, with its exciting emphasis on computational thinking, is under-funded. Good REU sites and programs have been terminated to allow new ones to start, and we were not shown evidence to support the belief that REU programs become self-sustaining. If these successful REU sites shut down when funding was shifted, the effect was to punish success.

Part C asked about forward-looking programs, and we see a plethora of opportunities – indeed, so many that even if CNS sees a substantial increase in funding in coming years, it can't possibly avoid making tough choices that may leave nationally important (even critical) topics underfunded.

Our COV does wish to voice some concerns:

- We cited the CRA and NAS studies above; both argue that increased funding to the systems research community could be beneficial, and both suggest that subsequent to the DARPA pullout, the overall (summed over all funding agencies) level of funding to this community dropped sharply. In addition, one important style of research, which could be called the "large-scale systems approach"—one or a small number of PIs, at a single institution, supporting a relatively large number of graduate students and staff—may not be well served by currently available government funding sources. Our COV took interest in this topic and attempted to evaluate the extent of such issues within the research community served by CNS.
    - Inadequate information on this topic was available for us to answer this question.
    - We feel that CNS should consider surveying its research constituency to quantify the effectiveness of its research programs in supporting computing and networking systems research, as perceived by the individuals actually doing the research. Doing this will not resolve questions about the financial impact of such events as the DARPA pullout or the 2008 financial crisis, but will clarify the current situation.
    - CNS has worked with many different programmatic structures in the general area over the years: ITR models, multi-institutional centers, smaller directly-funded efforts, GENI. Dialog with the research community could also clarify the question of which models work best for the people actually doing the work, while also yielding factual insight into the adequacy of overall funding levels today. In particular, we are curious to know how the overall level of overhead has evolved for CNS-supported researchers, as a function of the amount of research they are doing. Clearly, the goal should be to maximize research and minimize overhead, without sacrificing oversight and community-building.
- CNS should improve its mechanisms for tracking the impact both of new policies (such as once-annual submission deadlines) and also of the status of its larger programs, especially if cross-

funded (such as MRI, REU, GENI) or constituted as large centers with internal proposals, review processes and funding allocations (such as GENI, TRUST, etc).

- Policies such as one-per-year submission deadlines may be having unintended, harmful consequences. For example, new junior faculty members, who often start their positions in September, lack time to prepare proposals for a deadline that occurs more or less instantly, forcing them to wait a full year before submitting their first proposal to CNS. CNS should maintain a dialog with the impacted community about the costs and benefits of these new policies.

- CNS is doing an acceptable job of tracking managerial and other administrative issues created by cross-funded programs, but finds itself one level removed from the actual research decisions made by teams such as the GENI management office. Greater transparency of fine-grained research decisions would provide NSF (and future COVs) will know precisely what is being funded, what policies are used to select the funded work, what results are anticipated, and what is actually achieved.

- As CNS staff turnover occurs, CNS should institutionalize mechanisms for educating new staff about the status and accomplishments of complex, cross-cutting efforts so that knowledge acquired by the staff members who created a program will not be lost as that program is handed off to new generations of program officers at CNS.

- We were extremely impressed with the quality of work executed by program officers. These staff members have a uniquely broad perspective on CNS research that, if more accessible to the community, could increase the impact of CNS research. We make several observations on NSF process changes that could leverage this resource.

  - We are concerned that CNS program officers may lack time, support, or resources for functions beyond core proposal management. Program officers may spend excessive time on clerical work, due to, for instance, restrictive FastLane policies or lack of staff support (the "science assistants" apparently available in other NSF directorates). It is unclear to us whether CNS loads are genuinely higher than in other directorates, or whether other directorates have higher staffing levels. We do note that that statistically, CNS seems to be achieving a similar success rate to the rest of NSF. On the other hand, if this is being achieved by cutting budgets relentlessly, it is possible that the good success rates are being achieved at the cost of an unusually high effort level on the part of program managers.

  - ***At a minimum, we believe that CNS program officers must have funding to attend one or two top-ranked conferences per year.*** Not only does this afford a chance for the officers to keep a finger on the pulse of their research communities, it also permits them to network, to identify potential future reviewers and panelists, and to learn about hot emerging topics that may merit inclusion in new programs.

- We also believe that ***program officers must be given funding to undertake site visits for large, complex projects*** that use "internal" decision processes to make internal funding allocations, and during those site visits should conduct reviews of the decision process, the overall set of proposals received, and the review process used to select winning efforts.

- CNS walks a delicate tightrope reflecting tensions between focused programs and programs that welcome a wide breadth of ideas and out-of-the-box research.
  - CNS should actively maintain the balance one observes in (most) current programs. We say "most" because some clusters noted issues. For example, the review of NeTS revealed a program that had been broad and balanced in 2006-2007, but became smaller and narrowly focused in 2008. It is too early to evaluate the impact of this shift on the quality of the work being done.
  - Similarly, with respect to the major research instrumentation and academic infrastructure programs, it is important to recognize the inherent national interest in injecting cutting edge technologies into our research and educational programs (today, that might include very high speed networks, multicore platforms and virtualization, but this list will evolve rapidly as we move forward in time).
  - Proposals that offer innovative ideas for working with these cutting edge technologies should be encouraged, and panels should be made aware of the importance of achieving a high level of national sophistication with emerging commercial technologies.
  - This goal of increasing national sophistication with cutting edge technology must not eliminate funding for exciting but unusual ideas that innovate from outside of the box.
  - When attempting to foster community building, CNS must work with a light hand. The goal of community building must not interfere with the more primary goals of achieving the highest possible quality of scientific research, education and outreach.

Looking to the future, CNS enjoys an almost excessive number of opportunity areas, including such topics as multicore, virtualization, cloud computing, very high speed networks, and the implications of the new "network sciences" (the mathematical discipline driven by social networks, the web, and associated phenomena) for practical systems and communication networks. Our COV urges CNS to cooperate with industry and its research community to evaluate these and other promising opportunities, and the best ways to approach them. There are plenty of exciting waves of "disruptive" technologies, and CNS has a vital national role in simple ensuring that we extract the maximum value from each. Doing so may create entire new industries and revitalize the economy. Technology has historically been a huge engine for US economic growth and health, and in light of the staggering range of disruptive technologies suddenly reaching the market, there is every reason to expect that the same opportunities exist today.

SPECIFIC COV RECOMMENDATIONS

Primary Recommendations

1. CNS (or CISE as a whole) should develop a plan to evaluate impact of the various "proposal rate limiting" mechanisms created in 2008, to ensure that they are accomplishing the desired objectives while mitigating any unintended harm. This evaluation should include dialog with the research community impacted by the changed rules.

2. CNS should survey the community, asking: *Are researchers funded adequately, using efficient program structures, to do the best possible work with the lowest levels of overhead consistent with quality, oversight and community building?*

3. CNS should educate the reviewer community about ratings by providing NSF-wide statistics to all reviewers (permitting them to calibrate the scale used for reviewing)

4. CNS must also maintain program balance, engage the research community in establishing and periodically refocusing CNS programs, and take action to ensure that there will be broad community buy-in and participation in tactical decision making.

5. We commend CNS for its dialog with community, finding that:

    - Education workshops are highly effective. Town-hall meetings have also worked well, provided that the community showed enthusiasm and engagement.

    - Workshops have been less effective when perceived by community as "show and tell" and this type of event should therefore be avoided.

    - CNS should consider using research studies/workshops to identify focused investment opportunities and sketch out potential new programs.

6. Exploiting current technology changes is critical for national competitiveness

    - The nation must find ways to inject cutting-edge technologies into research institutions

    - CNS should incent those institutions to innovate on nationally critical new technologies. Potential examples might include fast networks, multicore, virtualization, cloud computing, and research on social networks and the associated science.

    - Injecting technologies can be "its own reward": proposals that offer smart ideas about using new technology should be recognized for doing so, not criticized for lacking completely new, out-of-the-box concepts.

7. The COV was pleased to see that proposals with excessive budgets relative to the work proposed were appropriately adjusted by POs. However, we also identified cases in which CNS seems to be stretching limited funds by negotiating budgets and timelines. Reducing budgets that were *not* excessive relative to the work proposed increases the success rate for high-quality proposals, but has serious impacts: reducing the scope of research projects and, over time, distorting proposals as PIs come to expect cuts and mold their proposals accordingly.

    - The best solution is to increase overall funding levels for CISE and CNS so that projects can be adequately funded without negatively impacting the number of projects funded.
    - A single-minded focus on success rates can incent POs to negotiate budgets even where the proposal budget had no fat to cut. This is highly undesirable.
    - As an effort is stretched in time, or cut beyond the threshold of effectiveness, an unreasonably high percentage of CNS dollars will be spent on non-research expense relative to dollars spend on research.
    - On the other hand, we do not suggest that budget negotiations aren't needed; we identified examples of proposals that clearly had excessive budgets relative to the work proposed. Budget negotiations should be used to correct such imbalances.

8. We noted some irregularity in the use and value assigned to "broader impact" statements. Some areas make effective use of these statements. In others, the concept of broader impact is less obvious. CNS may need to educate proposal writers, perhaps by sharing examples of particularly good broader impact statements used by others funded in the same programs.

9. NSF and CNS might benefit from a means of tracking programs using modern project management tools (allowing timeline analysis after the fact)

    - This includes tracking the impact of changed submission rules.
    - We are particularly concerned that the quality of information about "cross cutting" efforts and larger centers seems to be weak.

10. Program officers need more logistic resources, such as:

    - Program officers lack adequate personal travel resources.
    - In some (not all) clusters, reviewer incentives are needed. CNS program managers mentioned that the Westin Hotel has been unexpectedly helpful in this respect: reviewers seem enthusiastic about staying in what they perceive as a luxury hotel, yet the cost to NSF is the same.
    - More staff support would have a dramatic positive impact.

COV PROCESS RECOMMENDATIONS

Other Comments

1. Our group felt that it was given a difficult task, with inadequate time, and little opportunity to pre-plan and prepare.

    - The COV chair and cluster chairs were overwhelmed and lost time developing a plan for carrying out the COV study.

    - In contrast, NSF preparation and support was outstanding, and this was extremely helpful

2. Much could be done to make the COV process more effective in future iterations

    - Improved pre-visit planning would have been helpful (COV members arrived somewhat unprepared).

    - Cluster leaders should arrive a bit early (half day) to think through the task

    - Selection of a chair with prior COV experience would be useful

    - Our COV found that time spent in direct one-on-one dialog with program officers was far more useful than the broad program reviews that consumed almost half of our first day

    - Better data would help

        - We had no simple way to study panel composition, the full set of proposals considered by particular panels, or to review cross-panel statistics.

        - CNS could survey the research community using electronic means; doing so might shed light on many open questions.

        - Our COV did not have direct access to CNS databases and would have found this more useful and perhaps more effective than eJacket. Further, although it is useful to have determined a sample of proposals ahead of time, with specific representativeness characteristics, all proposals should be readily available, not just those in the sample. This would have saved a lot of time in the COV's analysis efforts.

        - Improved sampling methodology is desirable. The 2009 COV sample was performed by selecting every 20'th proposal from a name-sorted list of projects, and then by performing sub-sampling within efforts that were underrepresented in the initial selection. It was not clear to us that this process is adequately random. For example,

the initial selection clearly biases in favor of programs that received large numbers of proposals ; and given that the initial selection included renewals and increments as well as newly funded proposals, the initial selection had an unexpectedly large number of accepted proposals.

3. COV did not have access to information about workshops used by CNS to elicit new directions and opportunities from research communities, and to define priorities. We learned of them mostly through dialog with program officers.

   - Workshop documents should be well-publicized.
   - The COV was not shown any internal CNS documents articulating priorities for fostering new research directions.
   - It was difficult to navigate current and past solicitations, since, for example, programs (and their names) change very frequently.
   - COV had no “timelines” available to document the evolution of programs or relationships between them.

Committee of Visitors Members

Dr. Kenneth Birman, Chair
N.Rama Prof. of Computer Science
4119 Upson Hall
Cornell University
Ithaca, NY 14853
607-255-9199
ken@cs.cornell.edu

Dr. M. Brian Blake, Co-Chair
Computer Science
323 St. Marry's Hall
Georgetown University
Washington, DC 20057-1232
202-687-3084
Mb7@georgetown.edu

Cluster I. Computer Systems
Includes:

Computer Systems Research (CSR)
Distributed and Mobile Systems and Services (DMSS)
Parallel Systems, Computing and Execution (PSCE)

Cluster I. Chair
Dr. Fred Douglis
Currently changing positions; Previous address:
Research Staff Member
IBM T.J. Watson Res Center
19 Skyline DriveHawthorne, NY 10532
908-872-7335
Fd-nsf09@douglis.org

Dr. Remzi Arpaci-Dusseau
Computer Science Department
University of Wisconsin
1210 W. Daytona Street
Madison, WI 53706
608-263-7764
remzi@cs.wisc.edu

Dr. Clas Jacobson
United Technologies Research Ctr
Mail Stop 15
East Hartford, CT 06103-2607
860-610-7652
JacobsCA@utrc.utc.com

Dr. Aman Kansal
Microsoft Research
One Microsoft Way
Redmond, WA 98052
Kansal@microsoft.com

Dr. Edward W. Kohler
Computer Science
4531C Boelter Hall
University of California
Los Angeles, CA 90095-1596

Dr. Sanjay Shakkottai
Electrical & Computer Engineering
University of Texas
Austin, TX 78712
512-471-5376
shakkott@ece.utexas.edu

Dr. Jim Waldo
Distinguished Engineer
Sun Microsystems Labs
1 Network Dr, Bldg 2 ubur02-311
Burlington, MA 01003

NSF Resource Persons: Helen Gill
Michael Branicky
Krishna Kant
Anita La Salle

Cluster II. Network Systems
Includes:

Networking Technology and Systems (NeTS)
Networking Broadly Defined (NBD)
Wireless Networks (WN)
Future Internet Design (FIND)
Sensor Networks (SN)
Network Infrastructure (NI)

Cluster II. Chair

Dr. Joe Touch
Information Science Institute
University of Southern California
4676 Admiralty Way
Marina del Ray, CA 90292
310-448-9151

Dr. John Chuang
School of Information
102 South Hall, #4600
University of California
Berkeley, CA 94720-4600
510-642-7253
chuang@ischool.berkeley.edu

Dr. Bijan Jabbari
ECE Department, MS 1G5
George Mason University
Fairfax, VA 22030-4444
703-993-1618
bjabbari@gmu.edu

Dr. Philip Levis
Computer Science
358 Gates Hall
Stanford University
Stanford, CA 94305-9030
pal@cs.stanford.edu

Dr. Catherine P. Rosenberg
Electrical & Computer Engineering
University of Waterloo
200 University Avenue West
Waterloo, Ontario, Canada N2L 3G1
529-888-4510
cath@ece.uwaterloo.ca

Dr. Pravin P. Varaiya
Electrical Eng & Computer Science
271M Cory Hall
University of California
Berkeley, CA 94720
415-642-5270
varaiya@eecs.berkeley.edu

Dr. Walter Willinger
AT&T Labs – Research
Bldg 103, Room 229
180 Park Avenue
Florham Park, NJ 07932-0971
973-360-8419
walter@research.att.com

Dr. Rene L. Cruz
Department of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093-0407
cruz@ece.ucsd.edu
Phone: (858) 534-1426
Fax 1: (858) 822-1247

NSF Resource Persons: Darleen Fisher
Alhussein Abouzeid
Sajal Das
Victor Frost

III. Cyber Trust
Includes:

Computing Research Infrastructure (CRI)
Infrastructure Acquisition/Development (IAD)
Community Resource Development (CRD)
Planning Track (PT)

Cluster III. Chair

Dr. Steven Gribble
Computer Science & Engineering
University of Washington
578 Paul G. Allen Ctr, Box 352350
Seattle, WA 98195-2350
gribble@cs.washington.edu

Dr. George Davida
EE& CS
P.O. Box 784
University of Wisconsin
Milwaukee, WI 5320l
414-229-5192
davida@uwm.edu

Dr. Sonia A. Fahmy
Computer Science
Purdue University
305 N. University Street
West Lafayette, IN 47907-2107
765-494-6183
Fahmy@cs.purdue.edu

Dr. Petros Maniatis
Intel Research Berkeley
Intel Corporation
2150 Shattuck Avenue, Suite 1300
Berkeley, CA 94704-1345
510-495-3087
petros.maniatis@intel.com

Dr. R. Sekar
Department of Computer Science
State University of New York
Stony Brook, NY 11794
631-632-5758
sekar@cs.sunysb.edu

Dr. Scott Stoller
Computer Science
Stony Brook University
Stony Brook, NY 11794
631-632-16627
stoller@cs.sunysb.edu

Dr. Stewart Tansley
External Research & Program
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
425-707-5686
stansley@microsoft.com

NSF Resource Persons: Karl Levitt

IV. Research Infrastructure
Includes:

Computing Research Infrastructure (CRI)
Infrastructure Acquisition/Development (IAD)
Community Resource Development (CRD)
Planning Track (PT)

Cluster IV. Chair

Dr. Jorg Liebeherr
Electrical & Computer Eng

Dr. Melissa O'Neill
Computer Science
Harvey Mudd College

University of Toronto
10 King's College Road
Toronto, Ontario, M55, 3G4 Canada
416-946-3403
Jorge@comm.utoronto.ca

1250 North Dartmouth Avenue
Claremont, CA 91711-5980
909-607-9661
oneill@cs.hmc.edu

Dr. Dmitri Perkins
Ctr for Advanced Computer Studies
University of Louisiana
414 Mill Pond Drive
Lafayette, LA 70504
337-482-6732
perkins@cacs.louisiana.edu

Dr. Anand Tripathi
Department of Computer Science
University of Minnesota
200 Union Street SE, EECS 4-192
Minneapolis, MN 55455
612-625-9515
tripathi@cs.umn.edu

NSF Resource Persons: Anita LaSalle
Rita Rodriguez

Cluster V. Education and Workforce
Includes:

Broadening Participation in Computing (BPC)
CISE Pathways to Revitalized Undergraduate Computing Education (CPATH)
Research Experiences for Undergraduates (REU Sites)

Cluster V. Chair

Dr. Pamela B. Lawhead
Computer Science
373 Kinard Hall
University of Mississippi
University, MS 38655
662-915-5356
Lawhead@cs.olemiss.edu

Dr. Monica Brockmeyer
Computer Science
450 State Hall
Wayne State University
Detroit, MI 48202
313-577-0662
mab@cs.wayne.edu

Dr. Andrea W. Lawrence
Dept. of Computer Science
Spelman College
350 Spelman Lane, #362
Atlanta, GA 30314
404-270-5875
lawrence@spelman.edu

Dr. Gregory E. Monaco
Great Plains Network
Kansas State University
492 Bluemont
Manhattan, KS 66502
785-554-9657
greg@greatplains.net

NSF Resource Persons: Jan Cuny
Harriet Taylor
Joan Peckham

COV AGENDA

May 10 -- Evening Arrival

May 11 – All Visitors, Cluster Chairs, COV Chair

[Open Meeting]

8:00 - 8:30 AM	Morning Refreshments	Room 375, Stafford I
8:30 - 8:45 AM	Welcome/Charge to the CoV	Jeannette Wing, AD, CISE
8:45 - 10:00 AM	Division Overview, Purpose, Plan, Expectations	Ty Znati, DD, CNS
	Logistics, Program Data, Report Generation	Rajinder Khosla, DDD, CNS
	COI and Confidentiality Overview	Anita LaSalle, PD, CNS
10:00 – 10:15 AM	Break	
10:15 – 10:45 AM	CoV Sub-Group Assignments	Kenneth Birman, Chair M. Brian Blake, Co-Chair
10:45 – 12:00 PM	Introduction to CNS and its Clusters: 10 minutes for reports from NSF Resource Program Officers with Q/A	Cluster I – Helen Gill Cluster II – Darleen Fisher Cluster III – Karl Levitt Cluster IV – Anita LaSalle Cluster V – Jan Cuny and Harriet Taylor

[Closed Meetings]

12:00 - 1:00 PM Working Lunch: Jacket Review Procedure – eJacket Demo and Interactive Q&A Session

Cluster I	Computer Systems (CRS)	Room 375, Stafford I
Cluster II	Network Systems (NeTS)	Room 330, Stafford I
Cluster III	Cyber Trust (CT)	Room 370, Stafford I
Cluster IV	Computing Research Infras (CRI)	Room 1120, Stafford I
Cluster V	Education & Workforce (EWF)	Room 1150, Stafford I

1:00 – 3:30 PM	Review Data and Jackets: Clusters Devise Plan for producing Cluster Reports and begin reviewing COV materials	Room – same as in a.m.
3:30 -- 3:45 PM	Break (afternoon refreshments)	Room 375, Stafford I
3:45 – 4:45 PM	All Clusters reconvene for discussion and feedback session	Room 375, Stafford I
4:45 – 5:45 PM	Meet with Dr. Znati for Q&A and further information	Room 375, Stafford I
6:00 -- PM	Adjourn for Group Dinner	

Date – All Visitors, Cluster Chairs, COV Chair

(Closed Meetings)

8:30 - 8:45 AM	Morning Refreshments	Room 375, Stafford I
8:45 -- 10:45 AM	Continue to Review Data and Jackets	Clusters in same rooms as Wed. Stafford I
10:45 – 11:00 AM	Break	
11:00 – 12:00 PM	Continue to Review Data and Jackets	Clusters in same rooms as Wed. Stafford I
12:00 -- 1:30 PM	WORKING LUNCH All Clusters reconvene for discussion and feedback session	Room 375, Stafford I
1:30 – 3:00 PM	Prepare reports on Process and Outcomes	Clusters in same rooms as Wed. Stafford I
3:00 – 3:15 PM	Break	
3:15 - 4:15 PM	Cluster Meeting to prepare composite	Room 375, Stafford I
4:15 – 5:30 PM	Full CoV - Presentation of Cluster Reports and Discussion of Findings for Final Report Cluster Reports	Room 375, Stafford I
5:30 PM	Adjourn	

Date – Cluster Chairs, COV Chair

(Closed Meetings)

8:30 - 8:45 AM	Morning Refreshments	Room 390, Stafford I
8**:45 – 10:00 AM**	Full CoV - Discussion of Future Directions and Logistics of Report Completion	Room 390, Stafford I
10:00 - 10:30 AM	Executive Session with Assistant Director and Division Director	Room 390, Stafford I
10:30 – 1:00 PM	Cluster Chairs and COV Chair integrate Cluster Reports, outline Executive Summary, and outline Report to CISE Advisory Committee	Room 390 Stafford I
1:00 PM	Meeting Adjourn	

COV Coordination: Ty Znati, Rajinder Khosla, Gwen Owens

Cluster I Computer Systems (CRS)

FY 2009 REPORT TEMPLATE FOR
NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV:
Program/Cluster/Section:
Division:
Directorate:
Number of actions reviewed: **Awards:** **Declinations:** **Other:**
Total number of actions within Program/Cluster/Division during period under review: **Awards:** **Declinations:** **Other:**
Manner in which reviewed actions were selected:

PART A. INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

A.1 Questions about the quality and effectiveness of the program's use of merit review process.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE[1]
1. Are the review methods (for example, panel, ad hoc, site visits) appropriate? Comments: **Yes, but please see A.4.1 for comments on site visits.**	
2. Are both merit review criteria addressed In individual reviews? **Yes** In panel summaries? **Yes** In Program Officer review analyses? **Yes** Comments:	

3. Do the individual reviewers provide substantive comments to explain their assessment of the proposals? **Yes**

Comments: The committee found a surprisingly large fraction of sampled reviews offered useful, detailed, and substantive feedback to the proposers. One committee member described the sampled reviews as above the standard of some computer science conferences.

4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)? **Yes**

Comments: The committee was uniformly impressed by the extremely high quality of panel summaries and review analyses. CSR Program Officers are commended for the high quality proposer feedback they elicit from panels, and for the careful consideration of proposals, reviews, panel dynamics, and other constraints evident in review analyses.

5. Does the documentation in the jacket provide the rationale for the award/decline decision?

(Note: Documentation in jacket usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), program officer review analysis, and staff diary notes.) **Yes**

Comments: See comments above

6. Does the documentation to PI provide the rationale for the award/decline decision?

(Note: Documentation to PI usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), and, if not otherwise provided in the panel summary, an explanation from the program officer (written or telephoned with diary note in jacket) of the basis for a declination.)

Comments: **Yes**

7. Is the time to decision appropriate?

Note: Time to Decision --NSF Annual Performance Goal: **For 70 percent of proposals, inform applicants about funding decisions within six months of proposal receipt or deadline or target date, whichever is later**. The date of Division Director concurrence is used in determining the time to decision. Once the Division Director concurs, applicants may be informed that their proposals have been declined or recommended for funding. The NSF-wide goal of 70 percent recognizes that the time to decision is appropriately greater than six months for some programs or some individual proposals.

Yes

Comments:

8. Additional comments on the quality and effectiveness of the program's use of merit review process:

No additional comments

A.2 Questions concerning the selection of reviewers.

SELECTION OF REVIEWERS	YES , NO, DATA NOT AVAILABLE, or NOT APPLICABLE[2]
1. Did the program make use of reviewers having appropriate expertise and/or qualifications? Comments: **Yes**	
2. Did the program use reviewers balanced with respect to characteristics such as geography, type of institution, and underrepresented groups? **Yes** Note: Demographic data is self reported, with only about 25% of reviewers reporting this information. Comments: For underrepresented groups, CSR Program Officers have successfully tapped the admittedly limited pool of available reviewers. The committee commends the Program Officers on improving this aspect of reviewer balance since the last COV.	
3. Did the program recognize and resolve conflicts of interest when appropriate? **Yes** Comments:	

Additional comments on reviewer selection: No additional comments

A.3 Questions concerning the resulting portfolio of awards under review.

RESULTING PORTFOLIO OF AWARDS	APPROPRIATE, NOT APPROPRIATE[3], OR DATA NOT AVAILABLE
1. Overall quality of the research and/or education projects supported by the program. **Appropriate** Comments:	
2. Does the program portfolio promote the integration of research and education? **Appropriate** Comments: The committee applauds the increasing integration of undergraduates into research visible in sampled awards. We hope this continues.	

3. Are awards appropriate in size and duration for the scope of the projects?

Appropriate

Comments: The Program Officers have funded awards with an appropriately wide range of sizes and durations. The project scopes, after revision, also appear appropriate for their corresponding funding levels. The committee found in our sample obviously high-quality (Competitive and Recommended) proposals that were declined given the program's current funding levels. These proposals promised, at times, high impact and high likelihood of success, although they were (according to review analyses) less polished than other proposals. The committee also observed that in revised budgets, broader impact often received the brunt of the proposers' cuts. (Example proposal numbers: 0720857, 0615280.) It is clear that additional funding could be put to good use in the current proposal pool. The committee believes that increased CSR funding would lead to the creation of high impact research proposals.

4. Does the program portfolio have an appropriate balance of:

Innovative/potentially transformative projects?

Appropriate

Comments: The committee feels that the SGER program is a great source of funding for early-stage innovative and potentially-transformative projects, but one that isn't sufficiently broadly known in the CSR community. More advertising or education might help expose this useful program.

5. Does the program portfolio have an appropriate balance of:

Inter- and Multi- disciplinary projects?

Appropriate

Comments:

6. Does the program portfolio have an appropriate balance considering, for example, award size, single and multiple investigator awards, or other characteristics as appropriate for the program? **Appropriate**
7. Does the program portfolio have an appropriate balance of: Awards to new investigators? NOTE: A new investigator is an investigator who has not been a PI on a previously funded NSF grant. **Appropriate**
8. Does the program portfolio have an appropriate balance of: Geographical distribution of Principal Investigators? **Appropriate**
9. Does the program portfolio have an appropriate balance of: Institutionnel types? **Appropriate** :

10. Does the program portfolio have an appropriate balance:

Across disciplines and sub disciplines of the activity?

Appropriate

11. Does the program portfolio have appropriate participation of underrepresented groups?

Appropriate

Comments: The program portfolio demonstrates appropriate participation from underrepresented groups, given the number of applicants. Further attempts to increase the number of applicants from underrepresented groups may bear fruit. (Within the self-study sample, applicants and awards including underrepresented-group PIs: 2006, 52 awarded, 23 declined; 2007, 23 awarded, 35 declined; 2008, 18 awarded, 32 declined.)

12. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.

Appropriate

13. Additional comments on the quality of the projects or the balance of the portfolio:

The committee very much appreciated the breadth of the CSR call, including its focus on core discipline, the POs' technique of incrementally seeding potential new areas over multiple years, and reliance on community feedback for direction. The committee feels that the current programmatic focus on fewer, broader calls (with synchronized deadlines) is a strong improvement over previous, more numerous, and more focused calls, and should be maintained.

A.4 Management of the program under review.

1. Management of the program.

Comments: CSR Program Officers do work of remarkably high quality. Review summaries, panel summaries, and calls meet a high standard almost uniformly. The Program Officers are commended for their commitment.

The committee notes that the current workload on Program Officers, which forces them to spend most time on core proposal management functions, denies the computer science community access to their valuable expertise -- for example, in directing new research and in guidance for existing awards. The national interest would be served by giving POs the resources required to broaden their reach and their visibility within the community. Several specific concerns arose, in addition to a more general sense that the number of POs may be too low:

- Computer science research, particularly in systems areas, is defined by conferences and meetings, several of which happen each year. NSF Program Officer presence at conferences and meetings would greatly benefit the community. It would facilitate communication with PIs; it would allow POs to take the field's pulse and track funding priorities and upcoming opportunities; and by placing POs in more direct contact with research outcomes, it could let POs more proactively shift future funding opportunities. Cuts in travel budgets prevent this interaction, as do conflict-of-interest policies that make it difficult for POs to attend conferences with their own funds. Travel budgets should be raised.
- Program Officers may spend excessive time on clerical work, including work required by

restrictive policies in FastLane. Additional staff support -- perhaps from "science assistants" as available in other Directorates, perhaps from more clerical staff -- would give POs more time to focus on essential tasks, program overviews, and award review. FastLane should be loosened to allow POs to override business policies when appropriate (for instance, when a deceased PI has not approved a final report); current enforcement is too rigid.

2. Responsiveness of the program to emerging research and education opportunities.

Comments: The committee finds the CSR program appropriately responsive to emerging research and education opportunities. The balance within CSR of core discipline work and emerging work is good; both types of awards are extremely valuable. CSR's (and CISE's) switch to fewer, more broadly general calls, rather than many specific calls, is especially appreciated.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

Comments: Yearly updates to the CSR call have been effective. The committee appreciated the way new funding directions were "seeded" in the CSR call, some of which grew into new programs (CPS). This model works, especially in combination with ongoing core-discipline awards (which tend to be in topics that only NSF funds). The current balance between core discipline and areas of opportunity should be roughly maintained; the committee did not feel, for example, that the balance of resources should be shifted to NSF-predefined, prescriptive areas of opportunity.

4. Responsiveness of program to previous COV comments and recommendations.

Comments: The committee feels that the CSR program has been generally quite responsive to previous COV comments. The last COV requested more documentation of Program Officer funding decisions; as mentioned above, this committee found those decisions excellently well documented. The number of underrepresented-minority reviewers has distinctly risen. A field in proposal reviews for reviewer confidence is currently being implemented. The proposals we reviewed in general demonstrated reasonable effort in the integration of research and educational goals, perhaps in response to NSF education.

5. Additional comments on program management:

All comments appear above

PART B. RESULTS OF NSF INVESTMENTS

B. Please provide comments on the activity as it relates to NSF's Strategic Outcome Goals. Provide examples of outcomes ("highlights") as appropriate. Examples should reference the NSF award number, the Principal Investigator(s) names, and their institutions.

B.1 <u>OUTCOME GOAL for Discovery:</u> "*Foster research that will advance the frontier of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering."*

Comments:

CNS achieved its outcome goal of fostering research. The following are examples of success projects in this area funded by CNS:

1. NSF Award Numbers: 0447741

Award Title: CAREER: E-textile-based Wearable Computing for Sensing User Motions

PI: Thomas Martin,

Institution Name: Virginia Polytechnic Institute and State University

E-Textile Garments, Highlight ID: 13982, Version: AC/GPA

Researchers at the Virginia Tech are developing e-textile garments that can sense their own shapes and their wearer's activities. Many e-textile applications need to be aware of the positions of the sensors embedded in the textile relative to the wearer's body and relative to each other in order to function properly. Tracking the location of its sensing and processing elements lets the garment adapt to changes in how the user is wearing the garment, such as when the user rolls up the sleeves of a shirt or jacket. Continuously sensing the garment's shape makes it possible to determine the wearer's activity, enabling the e-textile to adapt to the user's current context. This proactive adaptation is an example and a key aspect of pervasive computing. The software services, sensors, and sensor algorithms developed in this project permit the design of e-textiles that can adapt to a wide range of applications without extensive tuning of the devices.

An Electronic-textile garment fabricated for context awareness. The insets show sensors and processing elements, all of which are connected together by an on-fabric network.

Applications for e-textile-based "wearable" computing include human computer interfaces, medicine, home care, manufacturing and industrial processes, entertainment, and emergency response. For example, garments are being designed that can detect whether the wearer is moving, determine characteristics of the wearer's gait, and even detect whether the wearer is falling. In cooperation with the Intel Corporation, the Virgina Tech team also has developed an e-textile rug that can track the position and motion of persons walking on the fabric.

One research team member, Dr. Thomas Martin, received the NSF Presidential Early Career Award for Scientists and Engineers (PECASE) for his leadership in the emerging area of e-textiles.

Primary Strategic Outcome Goal:

Discovery: Foster research that will advance the frontiers of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering. (AC/GPA selected)

Secondary Strategic Outcome Goals:

Research Infrastructure: Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools.

This project's research contributes to innovation and invention in the emerging area of e-textiles, and it can be expected to enable advances in other disciplines such as medical technology and health care delivery, monitoring and situation awareness, and human-assistive and augmentation technologies.

E-textiles research is inherently multidisciplinary, involving computer engineering, computer science, textile engineering, and garment design. In addition, this particular e-textiles project has the potential to impact a variety of fields, including home health care, human computer interfaces, and biomechanics.

2. NSF Award Numbers: 0509327

Award Title: CSR-EHS: Techniques for Assuring the Safety and Reliability of Physical Computing Systems and Applications to Medical Devices

PI: Insup Lee,

Institution Name: University of Pennsylvania

Techniques for assuring the safety and reliability of physical computing systems and applications to medical devices

Highlight ID: 16543, Version: AC/GPA

Medical devices are increasingly complex software-intensive systems. Many of these devices, such as infusion pumps, pacemakers, and ventilators, are critical for the safety and even life of patients. Existing technologies for assuring safety of these devices are not sufficient due to the complexity of the devices as well as due to the unpredictable and highly variable environments that these devices operate in. This insufficiency is seen in recalls of critical medical devices, even after their regulatory approval. In this project, Prof. Lee from the University of Pennsylvania, in close collaboration with his colleagues from the Hospital of the University of Pennsylvania and Massachusetts General Hospital, is pursuing new techniques for achieving high levels of assurance of safety-critical medical devices.

One of the major research avenues in this project targets pacemakers, a life-critical implantable device. Starting from the informal requirements for a pacemaker, provided by the vendor Boston Scientific, researchers are undertaking a model-driven development effort that will result in a working prototype of the pacemaker, supplemented by an assurance case that provides evidence that the developed prototype is safe. To date, the research team has developed a model of the digital controller of the pacemaker that supports all operational modes specified in the requirements, providing complete traceability from the requirements to the model elements. That is, for each model element such as a state transition, they identified a paragraph in the requirements document that justifies the model element. This mapping from requirements to the model forms an important part of the assurance case. The research team then proved, using the state-of-the-art verification tool UPPAAL, that the controller model satisfied some of the high-level properties listed in the requirements document, as well as other critical properties that were identified through the dialogue with Brian Larson, a developer at Boston Scientific. On-going efforts target generation of executable code for a particular embedded microprocessor and extending the assurance case to cover the generated code.

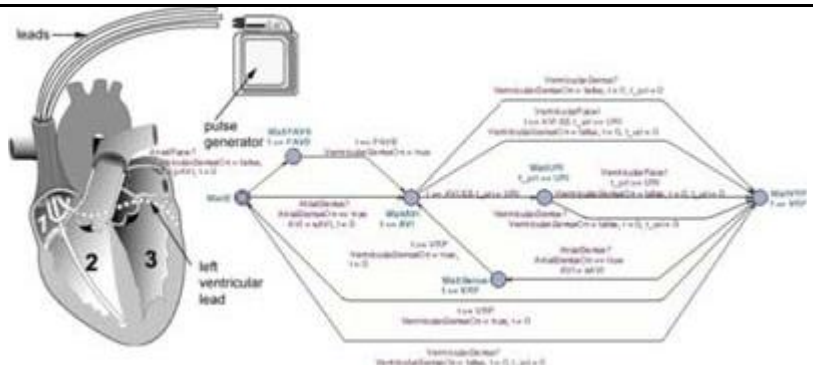

Formal modeling of cyber-physical systems, coupled with the careful construction of an assurance case, will lead to safer life-critical medical devices such as pacemakers.

The exercise of rigorous development and the construction of an assurance case for a real device will extend our understanding of validation and certification for safety- and life-critical cyber-physical devices. While vendors and government regulators are engaged in this practice on the regular basis, there is insufficient academic understanding of the practice and its research needs. Obtaining firsthand experience with the development and safety evaluation of safety-critical cyber-physical systems will let us target our research efforts to the more immediate and practical problems in the domain. More importantly, by presenting this experience in graduate and undergraduate courses on embedded systems design and validation, we will be able to train better engineers of the future.

B.2 <u>OUTCOME GOAL for Learning:</u> "*Cultivate a world-class, broadly inclusive science and engineering workforce, and expand the scientific literacy of all citizens.*"

Comments:

CNS achieved its outcome goal of cultivating its workforce and expanding scientific literacy. The following are examples of success projects in this area funded by CNS:

1.NSF Award: 0225610

Award Title: ITR: Foundations of Hybrid and Embedded Software Systems

PI: S. Sankar Sastry,

Institution Name: University of California-Berkeley

Highlight ID: 12297, Version: AC/GPA

Ideas: Hybrid systems compose continuous dynamics with discrete events. This project has developed an operational semantics for hybrid systems that gives a rigorous and precise meaning to models of hybrid systems by defining correctness criteria for software execution.

For one, discrete events may be simultaneous and yet causally related (e.g., instantaneous transfer of momentum). Also, simultaneous and causally related events are modeled by transient states, where exactly zero time is spent in the state.

Tools: A suite of tools provide support for multiple levels of computational hybrid systems. HyVisual (Hybrid System Visual Modeler) is an open-source software tool that uses a block-diagram representation of ordinary-differential-equations (ODEs) to define continuous and bubble-and-arc diagrams of finite state machines to define discrete behavior.

People: HyVisual, GReAT, Giotto and other tools are used at Berkeley and Vanderbilt in graduate classes on hybrid systems. Their visual syntax lends itself ideally to pedagogy, making it easier for instructors to illustrate subtle properties of hybrid systems. The open-source nature of the software encourages interaction with industry, which we will expect will take up the methods if not the software itself in next generation software tools for system design. Several textbooks are being developed by faculty at Vanderbilt and Berkeley on hybrid and embedded systems. Outreach summer programs sponsored at Berkeley and Vanderbilt prepare senior students for the graduate school research experience-especially women and underrepresented candidates.

2. NSF Award Numbers: 0540178

Award Title: DDDAS-TMRP: Collaborative Research: Adaptive Data-Driven Sensor Configuration, Modeling, and Deployment for Oil, Chemical, and Biological Contamination near Coastal Facilities

PI: Craig Douglas,

Institution Name: University of Kentucky Research Foundation

Intelligent Sensors Track Ocean Pollutants

Highlight ID: 16438, Version: AC/GPA

The Exxon Valdez oil spill, at its time the largest oil spill ever, no longer ranks among even the top 50 largest oil spills globally. Oil spills are still among the biggest potential threats to coastal water and water supplies. In response to these kinds of environmental threats, researchers at the University of Kentucky developed a data-driven pollution model that will also be useful in proposing countermeasures to general biochemical and chemical contaminants, including those that could be used in a terrorist attack through coastal waters. This project develops a new class of intelligent sensors and predictive computational models for pollutant tracking and identification that can dynamically adapt to new and unexpected conditions in remote water bodies.

This data-driven, real-time modeling methodology may mitigate damage to ecology, especially in areas where water supply depends on seawater desalination. It may also help inform remediation plans and early warning systems such as those that impact fishing and coastal

tourism industries, for example, forecasting calamities from algae blooms (red tide) or spilled chemicals.

This project uses a variable, light-wave sensor array integrated into an ocean observational system. This system may be superior to most near coastal ocean models, which are typically wind driven but not contamination transport driven. The multi-scale mathematical models and computer simulations developed by these researchers is both wind and transport-driven, dynamically injected with observed ocean data.

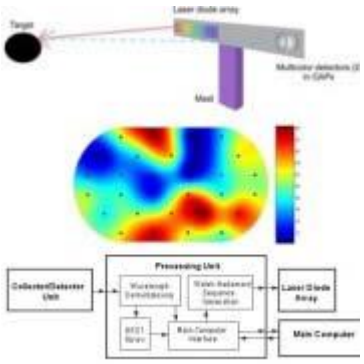

Solid-State Spectral Imager, spectrum image for analysis, and processing unit for intelligent sensing of pollutants.

These researchers are developing a symbiotic system between a new class of intelligent sensors and predictive computational models for pollutant tracking and identification that can dynamically adapt to new and unexpected conditions in remote water bodies. They are also educating a new generation of people who will be qualified to work in this environment which is close to real-time sensitive. This sensing system is able to go beyond sensing to prevention since the modeling connection enables small tank leaks to be tracked and repaired before major spills occur.

This project has one institution in an EPSCoR state (Kentucky). Graduate and undergraduates from Kentucky have participated in new courses developed as part of this project.

B.3 <u>OUTCOME GOAL for Research Infrastructure: "</u>*Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools."*

Comments:

CNS achieved its outcome goal of building the nation's research capability. The following are examples of success projects in this area funded by CNS:

1. NSF Award Numbers: 0205266

Award Title: ITR: Proactive Self-Tuning Systems for Ubiquitous Computing

PI: Daniel Siewiorek,

Institution Name: Carnegie-Mellon University

The testing and evaluation of wireless technologies is challenging because their performance depends strongly on the signal propagation environment, which, in turn, is determined by the physical environment. As a result, wireless testbeds are hard to isolate and control, and experiments are not repeatable. We developed a wireless network testbed based on signal propagation emulation that directly addresses these challenges. The testbed connects a diverse set of wireless devices, but instead of having the devices communicate through the ether, we capture their RF signals at their antenna ports and digitize the signals.

We then use FPGAs to emulate the effects of signal propagation in real time (e.g. large-scale attenuation, multi-path, small-scale fading, interference) and send the resulting signals to the wireless cards of the devices, after converting them into the RF domain. The wireless devices are placed in RF shielded boxes to ensure that they can only communicate through the emulator. The architecture can support all-to-all interference between all devices in the testbed. We have built a system that supports up to 15 devices operating in 2.4 GHz ISM band. The system as been deployed and made available to outside researchers over the Internet. We have used the system to demonstrate key properties of our approach. First, we can create and control very diverse experiments in software.

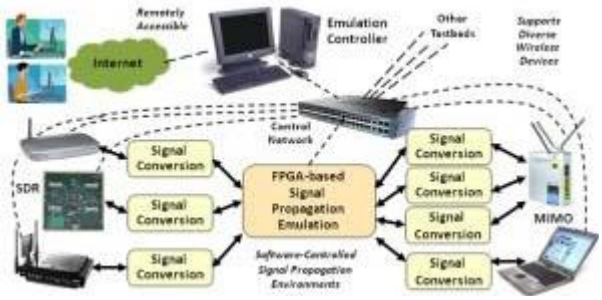

The emulator testbed uses signal propagation emulation to achieve ease of control and repeatability.

For example, we have used the emulator for experiments for vehicular networks, video streaming, automatic transmit rate selection, and ad hoc routing. Second, experiments at different layers in the system are supported. Example experiments include characterizing wireless 802.11 links (physical layer), automatically selecting transmit rate and transmit power (datalink layer), and video streaming (applications). Third, the wireless emulator testbed can support diverse devices (e.g. we recently started to experiment with software-radios besides 802.11) and it can be combined with wired testbed for hybrid wired-wireless experiments. Finally, experiments are fully repeatable at the physical layer. The emulator testbed is also a

useful tool for education: we developed an undergraduate course in wireless networking that uses the emulator for all assignments and projects.

Performing wireless experiments that are realistic, easy to control, and fully-repeatable is a very challenging problem. Many researchers are using testbeds "in the wild" - while they are realistic, they are hard to control (e.g. for mobile devices) and results are not repeatable and are notoriously hard to interpret. Simulation is a poor substitute. While experiments are repeatable and easy to control, there is significant concern about realism - today's wireless devices are complex and it is difficult to simulate (re-implement) all relevant aspects of the system. The wireless emulator combines the benefits of both of these approaches. We have been able to show that the combination of realism, control, and repeatability is not only useful for evaluating wireless technologies, it also helps in understanding and characterizing the behavior of wireless systems and in optimizing their performance.

2. NSF Award Numbers: 0540160

Award Title: DDDAS-TMRP: Dynamic, Simulation-Based Management of Surface Transportation Systems

PI: Richard Fujimoto,

Institution Name: GA Tech Research Corporation - GA Institute of Technology

New Computer Models Help Ease Traffic Congestion

Highlight ID: 14600, Version: AC/GPA

Traffic delays in transportation systems cost the nation tens of billions of dollars annually in lost productivity, wasted fuel, and the consequences of traffic pollution. In times of crisis, congestion and delays can result in lost lives. The goal of this project is to ease traffic flow, particularly during unplanned events, by having traffic decision-support systems take advantage of accurate models of traffic phenomenon. This project’s research applies the foundations of dynamic data driven applications systems to the generation of ad hoc distributed traffic simulations. These simulations model collections of autonomous vehicles dynamically interacting with each other. The simulated vehicles are controlled by real-time data applied to dynamic models that seek to predict the occurrence of real scenarios based on actual traffic conditions.

Aerial view of urban region modeled by the DDDAS traffic simulator and predictor.

This project is an example of how the widespread deployment of sensors, computers, and communications in vehicles and roadways can be leveraged to improve safety. To effectively exploit the wealth of real-time data and information that is becoming available, researchers are developing methods to selectively capture, process, and incorporate relevant data into simulation models and other decision-support tools and use predictive models both in normal day-to-day operations and in the presence of major disruptive events. Not only is the project developing simulation models capable of incorporating dynamically monitored data, it also addresses technical challenges of data distribution and synchronization.

In addition to its research components, the project is using realistic scenarios and test data to validate the approaches being used. Traffic simulations are being coded to model an area near Georgia Tech in midtown Atlanta, GA, with an evacuation scenario that focuses on traffic leaving the area via bridges that are typically bottlenecks to traffic flow.

PART C. OTHER TOPICS

C.1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

While we believe there were issues of focus as of the time of the previous COV, we believe the current broad focus is appropriate and allows the coverage area to be determined by the research community.

We identified one possible gap that is not specific to CSR but bears consideration. NSF sponsors the Cluster Exploratory (CLUE) program, which makes IBM/Google cloud computing clusters available to a number of universities. At present, CNS is participating in the Data Intensive Computing cross-cut, with an investment in excess of $3M. However, given the exceptional level of industry activity in cloud computing and the desire to broaden academic involvement both in education and research activity in this area, we encourage CNS to further investigate such cross-cutting activities targeting innovative research in cloud computing and virtualization, and also innovations in teaching students to work with cloud computing platforms and infrastructures.

C.2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

Past cuts in funding have made it difficult for NSF to attain the desired level of project funding and the full portfolio of research that would be desirable. Specifically, the need to fund applicants at lower levels tends to result in a reduction of the broader impacts those proposals can achieve.

C.3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

Issues with limited staffing and travel allowances are critical problems that need to be addressed. Please refer to section A.4.1 for detailed recommendations.

C.4. Please provide comments on any other issues the COV feels are relevant.

The committee appreciates instances of CSR's advocacy efforts for new technical directions, particularly to industry and academia. The committee observes that these advocacy efforts can result in increased proposal load. Without accompanying buy-in from NSF leadership, including, potentially, funding increases, the result can be high workload on Program Officers, reduced funding per proposal, and incomplete realization of the POs' strategic vision. We have no specific recommendations.

C.5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

The "ad hoc" meetings with program officers were extremely useful parts of our meetings. These interviews should be regular parts of the agenda and take place both early in the meeting (detailed discussions over overall research agendas, methodologies, and issues) and later (to discuss specific questions that arise).

We found the sampling method used for generating cluster-specific jackets to be minimally documented and possibly subjected to statistical issues. We recommend that the sampling technique be formalized and standardized.

We would like to have focused on full proposals, with the evaluation of certain other types (continuation, workshops, etc.) treated separately.

The self-study and previous COV reports should be required reading in advance of the meeting, along with documentation of the charge to the COV. This would allow attendees to prepare better and spend less time on-site coming up to speed.

Cluster II Network Systems (NeTS)

)

The table below should be completed by program staff.

Date of COV:
Program/Cluster/Section:
Division:
Directorate:
Number of actions reviewed: **Awards:** **Declinations:** **Other:**
Total number of actions within Program/Cluster/Division during period under review: **Awards:** **Declinations:** **Other:**
Manner in which reviewed actions were selected:

PART A. INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

A.1 Questions about the quality and effectiveness of the program's use of merit review process.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or

	NOT APPLICABLE[1]
1. Are the review methods (for example, panel, ad hoc, site visits) appropriate? Yes. Comments: Because of a limited travel budget and heavy workload, program managers are less able to make site visits, which can be a valuable component of reviewing larger projects. Panel reviews are used almost exclusively, compared to ad hoc reviews; we do not see this as an issue.	
2. Are both merit review criteria addressed In individual reviews? Yes In panel summaries? Yes In Program Officer review analyses? Yes Comments: In 2003-5, the percentage of reviews that did not address both criteria dropped from 10% to 5%. In 2006-8, the percentage rose back to 10%. In general, intellectual merit review is well addressed. Panel summaries tended to consist of merged text of individual reviews, and are an opportunity to provide context from the panel discussion where possible. Program Officer reviews tended to be much more constructive and provided a valuable synthesis of information. Broader impact merit review would benefit from clarification. Reviews tended to summarize responses in the proposals.	

3. Do the individual reviewers provide substantive comments to explain their assessment of the proposals? Yes.

Comments:

In 2003-5, the portion of insufficient reviews ranged from 4-10%. In 2006-8, the portion was 10-16%. This suggests that the depth of reviews is degrading over time.

We reviewed the 2006-8 report and noted that it recommended instituting a system to acknowledge and reward good reviewers, and the NSF said it would explore this idea. It is not clear what actions the NSF has taken to explore this issue. Evidence suggests that the issue is growing in severity. The NSF may want to revisit how to address it.

4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)?

Generally, yes.

Comments:

The panel summaries are an opportunity to explain a rationale for consensus, in addition to the rationale for the overall decision. Summaries could include more information on the panel's decision making process, rather than focusing on justifying the individual decision of a proposal.

5. Does the documentation in the jacket provide the rationale for the award/decline decision?

(Note: Documentation in jacket usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), program officer review analysis, and staff diary notes.)

Yes.

Comments:

In general, the program officer review analysis often provided the most thorough and constructive information on the rationale for a decision. We were unanimously impressed with and appreciative of the effort of the program officers in providing this insight and feedback.

6. Does the documentation to PI provide the rationale for the award/decline decision?

(Note: Documentation to PI usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), and, if not otherwise provided in the panel summary, an explanation from the program officer (written or telephoned with diary note in jacket) of the basis for a declination.) Yes.

Comments:

It is very important to keep the documentation quality as high as it has been. Constructive comments, as are common today, are very helpful. This is especially true for rejected proposals by young PIs. Again, the most useful information often was provided by the program officer's summary.

7. Is the time to decision appropriate?

Note: Time to Decision --NSF Annual Performance Goal: **For 70 percent of proposals, inform applicants about funding decisions within six months of proposal receipt or deadline or target date, whichever is later**. The date of Division Director concurrence is used in determining the time to decision. Once the Division Director concurs, applicants may be informed that their proposals have been declined or recommended for funding. The NSF-wide goal of 70 percent recognizes that the time to decision is appropriately greater than six months for some programs or some individual proposals. Yes.

Comments:

All programs had prompt responses with at least 71% of proposals, and in many cases there were timely responses for over 90% of proposals.

8. Additional comments on the quality and effectiveness of the program's use of merit review process:

Panels are doing a very thorough job. Although there is some evidence that not all members of every panel are fully prepared when the panel convenes, overall the quality of reviews and discussion at the conclusion of the panel are very high and helpful to proposal writers.

Clearer statements of expectation to panel members may result in deeper and more effective reviews. We note that prior COVs (Section 2.5 of the 2003 COV report, echoed in the 2006 COV report) have made concrete recommendations on how to improve reviews, yet there have not been significant changes in review forms or preparation. There is some discussion within the networking and systems community on peer and merit review: the WOWCS workshop at NSDI 2008 was held to discuss this. The NSF may benefit from reviewing these ideas to make its panel review process even more thorough and valuable.

This is especially true for the broader impact merit criterion. Intellectual merit is handled quite effectively, but the panel felt the broader impact statements and evaluation in regular technical proposals (i.e., not educational focus or supplements) gave the impression of cursory treatment by both the PIs and reviewers. Clearer expectations in this dimension of merit review would address this issue.

A.2 Questions concerning the selection of reviewers.

SELECTION OF REVIEWERS	**YES , NO, DATA NOT AVAILABLE, or NOT APPLICABLE[2]**
1. Did the program make use of reviewers having appropriate expertise and/or qualifications? Yes. Comments: Although program managers are able to assemble well-qualified panels from a broad area of backgrounds and research areas, doing so is labor intensive. This is in part due to the constraints placed on panel formation, such as the limited time available to gather the panel, and the COI rules. The current trend towards a single deadline of a single program may be amplifying these effects. It may be worthwhile to explore ways to increase the pool of eligible reviewers; in the past, this has sometimes been done by splitting panels by project size	

(small/medium/large) and/or by subtopic (wireless, sensors, NBD).

2. Did the program use reviewers balanced with respect to characteristics such as geography, type of institution, and underrepresented groups? Yes.

Note: Demographic data is self reported, with only about 25% of reviewers reporting this information.

Comments:

Computer science and engineering have historically had low representation among underrepresented groups, with systems and networking being some of the most troubled areas. The degree to which NSF panels include members of underrepresented groups is correspondingly laudable and excellent. Program managers are also doing a good job attracting reviewers from industry as well as from a mix of universities and research labs.

3. Did the program recognize and resolve conflicts of interest when appropriate? Yes.

Comments:

Based on our review of the COI guidelines and our own personal experiences, we have no concerns with the review process. The NSF process represents a high standard admired throughout the world in this regard.

We were particularly impressed that program officers do not permit panelists to engage in dialog with the PI of a pending proposals even for the purpose of eliciting clarifying information. To insure that as broad a perspective as possible can be provided to panelists, we recommend that proposals include supporting community letters, and encourage the NSF to follow up with the community at large (including applicants whose proposals were turned down) to guage the public perception and impact of NSF-funded work in this area.

4. Additional comments on reviewer selection:

A.3 Questions concerning the resulting portfolio of awards under review.

RESULTING PORTFOLIO OF AWARDS	**APPROPRIATE,** **NOT APPROPRIATE[3],** **OR DATA NOT AVAILABLE**
3. Overall quality of the research and/or education projects supported by the program. Appropriate.	

4. Does the program portfolio promote the integration of research and education?

Appropriate.

3. Are awards appropriate in size and duration for the scope of the projects?

Not appropriate.

Comments:

Limited funding and an increase in proposals have forced NSF to reduce grant sizes and durations in the areas studied by this cluster, at times to levels below what proposals seek to achieve. This has reduced the scope of research that PIs can tackle, and is rapidly approaching critical quantization points, e.g., managing funding for fractions of students is difficult. Given the constraints placed upon program managers, the COV believes the NSF is doing an excellent and exemplary job, but should continue to try to avoid fractional support for students. Further, the NSF should be careful to avoid triggering a gaming effect, in which PIs propose levels in the expectation of particular reductions. Reductions should be used only sparingly; it may be more effective (and avoid gaming) to reduce grant size in the request for proposals (announcing lower limits or targets) rather than via negotiations during the award phase.

4. Does the program portfolio have an appropriate balance of:

Innovative/potentially transformative projects?

Appropriate.

Comments:

The large commitment to SGER grants demonstrates the NSF's commitment to innovative

and speculative work. Furthermore, history has shown the NSF has funded transformational research that has, even in a short period, changed computing. Examples include the IRIS ITR grant, which laid the foundation for DHTs in the Internet today, cross-layer network design, and further examples that we present in Section B.

In particular, NeTS has a good track record of encouraging innovative/potentially transformative projects by its very design. During the 2006-2007 period NeTS included a "Networking Broadly Defined" category, which served as an incubator for new ideas that did not fit under the targeted directions of other sub-categories. In addition, the FIND program is specifically focused on encouraging such projects.

5. Does the program portfolio have an appropriate balance of:

Inter- and Multi- disciplinary projects?

Appropriate.

6. Does the program portfolio have an appropriate balance considering, for example, award size, single and multiple investigator awards, or other characteristics as appropriate for the program? Appropriate.

Comments:

With the future influx of funding, particular care should be given to new and young investigators, specifically the CAREER award. Current CAREER awards, due to inflation and increase in tuition, are unable to effectively support a new faculty member. For example, they cannot fully support a student for 5 years; this is another example of triggering the quantization effect noted earlier. The NeTS COV suggests that the NSF examine the original mission of the CAREER award and examine how it can return to its original intent. We do not recommend reducing the number of CAREER awards. We reiterate the importance of providing sufficient funding to CAREER awards without compromising the number of recipients each year.

7. Does the program portfolio have an appropriate balance of:

Awards to new investigators? Appropriate.

NOTE: A new investigator is an investigator who has not been a PI on a previously funded NSF grant.

Comments:

Within the NeTS cluster, the portion of grants including new investigators was 14% in FY06; 30% in FY07 and 22% in FY08. These are impressive rates, and may be particularly indicative of the ability of subprograms to successfully attract new investigators.

8. Does the program portfolio have an appropriate balance of:

Geographical distribution of Principal Investigators? Appropriate.

9. Does the program portfolio have an appropriate balance of:

Institution types? Appropriate.

10. Does the program portfolio have an appropriate balance:

Across disciplines and sub disciplines of the activity?

Appropriate.

Comments: We refer to Section C, where we applaud how NeTS has enabled deep focus sub-areas while simultaneously supporting breadth through general calls.

11. Does the program portfolio have appropriate participation of underrepresented groups?

Comments:

Appropriate.

Comments:

It is clear that the program managers are making an excellent effort, and should be commended for their commitment to this important issue.

12. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.

Appropriate.

Comments:

FIND in particular, and NetSE as currently planned, is in reaction to and in support of national priorities to reexamine the architecture of the Internet without the constraint of assuming particular incremental deployment strategies. Such “out of the box” programs may support the agency mission as well, by fostering new communities and concepts.

13. Additional comments on the quality of the projects or the balance of the portfolio:

None.

A.4 Management of the program under review.

1. Management of the program.

While there has been a great increase in the number of proposals and workload, the staffing level within CNS has not increased. 2006 saw a record number of proposals to CNS: 2,348. Despite this increase, the CNS staff has been able to manage its program very effectively and remain accessible to visitors, but we are concerned of the effect on CNS staff and its ability to support its PIs through site visits and other interactions that can be critical to creating and supporting a thriving research community.

Members of the COV found it difficult to gauge some aspects of management due to a difficulty to get information on PI meetings and grant reports. Greater visibility into management of programs between funding and completion would help the COV provide more feedback on this issue. FIND is a notable exception, because of its web pages that include meeting records across multiple years and panels in a single location

Comments: None.

2. Responsiveness of the program to emerging research and education opportunities.

Focus areas in the 2006-8 timeframe explored several emerging research areas that have proven to be important, such as wireless networks, embedded sensing networks, and network architecture.

Comments: None.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

The ratio between long-lived and shorter focus areas, providing a mix of stability and variation, was very effective. For example, NOSS had several years of focus, building a community of wireless sensor network researchers, while simultaneously ProWIN has seeded a community of

software radio research, extending into network coding and the physical layer.

Comments: None.

4. Responsiveness of program to previous COV comments and recommendations.

There are several recurring issues in prior COVs that have, to our best understanding, not been addressed. First and foremost among them is the ratio of staffing to workload. Current CNS staff handles a record number of proposals while simultaneously kicking off completely novel and labor-intensive programs, such as GENI. Simultaneously, a lack of support leads CNS staff members to spend time on handling their own IT problems. Despite this tremendous increase in work, staffing levels remain almost equal to their 2003 levels.

This overload can and will have broad-reaching negative effects to the quality of management that CNS can provide. For example, program managers are less able to engage the community in site visits, PI meetings, participating as a catalyst at conferences and other research meetings, and other valuable activities. This may also undermine the ability of CNS to attract excellent program managers in the future.

Comments:

5. Additional comments on program management:

None.

PART B. RESULTS OF NSF INVESTMENTS

B. Please provide comments on the activity as it relates to NSF's Strategic Outcome Goals. Provide examples of outcomes ("highlights") as appropriate. Examples should reference the NSF award number, the Principal Investigator(s) names, and their institutions.

B.1 OUTCOME GOAL for Discovery: "*Foster research that will advance the frontier of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering.*"

Comments:

NeTS use of targeted focus areas has enabled transformational research in several networking subfields. For example, the FIND program examines how to completely transform the current Internet to a new network with very different capabilities (James Sterbenz, Kansas University, 0626949). Similarly, NSF's funding of wireless research through WN and ProWIN has in part been responsible for breakthroughs in network coding and packet recovery (Hari Balakrishnan, MIT, 0721702). Finally, funding of embedded networks through NOSS has led to transformational approaches to low-power networking (David Culler, Berkeley, 0435454).

In addition to this transformative research, NeTS has also funded a great deal of foundational and deep work. For example, within NBD, there is a project that examines how to make the current Internet more energy efficient (Kenneth Christensen, University of South Florida, 0520081). Within FIND, there are similar projects that look at energy efficiency in terms of network architecture (Mark Allman, ICSI, 0721933). Finally, the foundational work on the capacity of multihop wireless networks, which has created a whole branch of research, was funded in part by the NSF (PR Kumar, Illinois, 0519535).

Other grants span both, investigating current problems with the potential to transform how they are tackled, such as market economics of wireless spectrum management (Walrand, Berkeley, 0627161).

B.2 OUTCOME GOAL for Learning: "*Cultivate a world-class, broadly inclusive science and engineering workforce, and expand the scientific literacy of all citizens.*"

The Network Cluster community has supported an active and inclusive workforce, and to expand the literacy of the community at large by providing numerous open software components, many of which are integrated into widely distributed systems, such as Linux, FreeBSD, as well as influencing proprietary platforms such as Windows, MacOS, and router and switch software. Some of these components are provided as software as a stand-alone package

or integrated into existing syste), whereas others are provided as services (e.g., some services deployed within PlanetLab). These are in addition to the infrastructure components that include facilities, instrumentation, and experimental tools noted in B.3.

Comments: None.

B.3 <u>OUTCOME GOAL for Research Infrastructure: "</u>*Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools."*

Understanding that networking research requires new infrastructure in the form of testbeds and experimental tools, NeTS has funded many successful projects that provide such resources. The Kansei initiative, which incorporates a federation of wireless testbeds into the GENI project (Anish Aurora, Ohio State, GENI project number 1633). Many wireless researchers use Kansei heavily. At the WINLAB in Rutgers, the Orbit Radio Grid (Dipankar Raychaudhuri, award 0725053) is another example of useful networking facilities, as is Tutornet at USC (Ramesh Govindan, 0520235). Many researchers use these facilities, which provide both experimental testbeds as well as repeatable experimental methodologies.

WARP (Wireless open Access Research Platform) provides a powerful infrastructure for novel wireless designs (Edward Knightly, 0224458).

PlanetLab (Larry Peterson, Princeton University, 0335314) has enabled a new generation of distributed systems and networked systems research by providing a global, highly heterogeneous testbed to researchers worldwide.

The deep and broad investment in infrastructure and facilities has been a critical enabler to an impressive range of networking research.

Based on our personal experiences, we believe that the maintainers of these resources actively try to make them available to the wider community: a survey of the community on their use and perceived accessibility may be helpful. In some cases, such as PlanetLab, this is clear. With smaller efforts, such as Tutornet, perhaps coordination through the NSF to advertise their availability could better leverage these critical investments.

We acknowledge that there is a tremendous expectation of how GENI may open new doors to future networking research.

Comments: None.

PART C. OTHER TOPICS

C.1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

NeTS is currently stated very broadly, such that there are no clear gaps. Instead, the current (2009) description may be too broad. The panel felt that previous strategies involving a set of targeted sub-areas (FIND, ProWIN, WN, etc.) that varied in a staggered fashion worked well when complemented by the NBD (Networks Broadly Defined) sub-area. During the review period, this continued through 2007, where areas (in general) focused on technologies. In 2008, the sub-areas focused on approaches; we felt this was a reasonable alternative approach worth exploring, but were concerned that the "other" category represented by NBD had been omitted at that time. The current (2009) description appears very broad, and may lack the community-building impact of previous targeted programs.

Overall, we felt that the balance of targeted areas together with a catch-all open area was the better approach, because the target areas help focus community efforts and leverage common interests and resources, whereas the open area serves as an unbiased incubator for projects that may lead to future target areas. We encourage the NSF to use this combined, balanced approach. NetSE helps contribute to this approach, providing a cross-area venue for targeted efforts in Network Science. When such a large investment is made in programs with a lack of focus areas, it would be useful to closely monitor those programs to determine whether introducing sub-areas would be of use.

C.2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

We felt that the NeTS program was generally successful in meeting its goals and objectives. This is especially true of FIND, intended to foster “out of the box” approaches, and which clearly has engaged new communities as a result.

We were concerned about the balance of intellectual merit goals vs. broader impact goals, that the broader impact could be misinterpreted as being handled in a cursory fashion. As noted elsewhere in this report, there are numerous ways in which we felt projects had broader impact that could be noted more prominently, e.g., in providing not just tools and instrumentation, or publishing results, but in contributing to an evolving software infrastructure that is just as critical to the educational and workforce support.

C.3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

We felt it would be useful to provide guidance to both PIs and reviewers on the variety of ways in which projects can have broader impact, to encourage them to address this impact in a more useful, individual, and specific manner.

C.4. Please provide comments on any other issues the COV feels are relevant.

None.

C.5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

Although we received ample information on previous COV reports and NSF responses, it would be useful to get direct information on specific actions that the NSF ultimately took to address the responses.

The COV could also provide more useful feedback on program management if we had more exposure to the process between project initiation and project completion. In particular, it would be useful to have information, organized by program (not just by panel), that collated project nuggets, project reports, and PI meeting info. For PI meetings, it would be useful to have not only the agenda of the meeting, but also a brief (one page) report on the result of the meeting, e.g., describing notable aspects of the meeting structure, observations or issues raised at the meeting, and an assessment of the utility and impact of the meeting.

Cluster III Cyber Trust (CT)

FY 2009 REPORT TEMPLATE FOR

NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV: 5/11/2009-5/13/2009
Program/Cluster/Section: Cyber Trust
Division: CNS
Directorate: CISE
Number of actions reviewed: **Awards:** **Declinations:** **Other:**
Total number of actions within Program/Cluster/Division during period under review: **Awards:** **Declinations:** **Other:**
Manner in which reviewed actions were selected:

PART A. INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

A.1 Questions about the quality and effectiveness of the program's use of merit review process.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE[1]
1. Are the review methods (for example, panel, ad hoc, site visits) appropriate? Comments: Yes. We found that panels are uniformly effective in recommending high-quality proposals to fund, and that site visits are useful both while selecting center-scale proposals to fund and as a mechanism for providing review and guidance to in-progress, large or center-scale grants. The self-study report notes that ~90% of proposals are reviewed exclusively by panels, which seems both appropriate and desirable: this allows the majority of proposals to benefit from uniformity and consistency of evaluation, while tending to reserve email reviews as a secondary measure for when appropriate panels (or perhaps panelists) cannot be found.	
2. Are both merit review criteria addressed In individual reviews? In panel summaries? In Program Officer review analyses? Comments: Yes. Both criteria are addressed in all three cases. We noted that intellectual merit tends to receive more attention, and consequently more substantive review, than broader impacts, but this is reflective of the content in the proposals themselves, and seems appropriate.	

[1] If “Not Applicable” please explain why in the “Comments” section.

3. Do the individual reviewers provide substantive comments to explain their assessment of the proposals? Comments: Yes. In most cases, reviewers provided substantive comments. We did find that some reviews were inadequately thorough, but this is the fault of the reviewers who wrote them and seems like an unavoidable issue. However, our analysis confirmed that any given proposal that we examined received multiple high quality reviews.
4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)? Comments: Yes. We found that most panel summaries are informative. A minority of summaries tended to mirror a dominant review, rather than reflecting the full texture and substance of what is, in our collective experience, lively discussion that occurs in panels. This may be a matter of time pressure, or perhaps a scarcity of resources available to panels. We did note that high quality proposals tended to receive consistently thorough reviews and panel summaries. On the few proposals in our sample where a dissenting review was present, the panel summary contained adequate justification to counter the dissenting reviewer's arguments
5. Does the documentation in the jacket provide the rationale for the award/decline decision? (Note: Documentation in jacket usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), program officer review analysis, and staff diary notes.) Comments: Yes. The program managers provided uniformly high quality and detailed analysis of decisions within the jackets we analyzed.

6. Does the documentation to PI provide the rationale for the award/decline decision?

(Note: Documentation to PI usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), and, if not otherwise provided in the panel summary, an explanation from the program officer (written or telephoned with diary note in jacket) of the basis for a declination.)

Comments:

Yes.

7. Is the time to decision appropriate?

Note: Time to Decision --NSF Annual Performance Goal: **For 70 percent of proposals, inform applicants about funding decisions within six months of proposal receipt or deadline or target date, whichever is later**. The date of Division Director concurrence is used in determining the time to decision. Once the Division Director concurs, applicants may be informed that their proposals have been declined or recommended for funding. The NSF-wide goal of 70 percent recognizes that the time to decision is appropriately greater than six months for some programs or some individual proposals.

Comments:

Yes, the CT cluster met this dwell time performance goal. The COV was curious about the justification behind the "70% over six months" goal, but suspected this is an appropriate goal.

8. Additional comments on the quality and effectiveness of the program's use of merit review process:

Overall, the merit review process is of high quality, and is very effective at identifying and selecting meritorious proposals for funding. To reiterate an earlier point, though it is perhaps inevitable that some reviews and some panel summaries lack appropriate thoroughness, the majority were of high quality, and we found that the overall process was robust in dealing with

the minority of lower quality reviews and summaries.

A.2 Questions concerning the selection of reviewers.

SELECTION OF REVIEWERS	YES , NO, DATA NOT AVAILABLE, or NOT APPLICABLE[2]
1. Did the program make use of reviewers having appropriate expertise and/or qualifications? Comments: Yes. Reviewers had appropriate expertise and qualifications, and the COV was particularly impressed with the quality of reviewers that were assembled to gauge large and center-scale proposals.	
2. Did the program use reviewers balanced with respect to characteristics such as geography, type of institution, and underrepresented groups? Note: Demographic data is self reported, with only about 25% of reviewers reporting this information. Comments: Yes. The self-study data confirms that there is a balance of reviewers across these characteristics: the balance is commensurate with what is present in our research community at large.	
3. Did the program recognize and resolve conflicts of interest when appropriate? Comments: Yes. The process by which COIs are reported and handled is very thorough and professional. No quantitative data on the implementation of the COI process during panels was available to the COV, such as occurrences of conflicts of interest and their resolution. However, based on COV members' experience, situations in which COI issues arise are handled in an appropriate and transparent manner.	

Additional comments on reviewer selection: The COV commends the CT program on the quality of its reviewers and reviews. We recognized that it is difficult to find a consistent supply of qualified reviewers, but in spite of these constraints, panels and individual proposal reviews were rigorous, and the funding recommendations made are of high quality. The COV was curious whether it would be feasible to gather data on reviewer invitations and

declinations, and to use this to infer or track issues related to reviewer fatigue or pool exhaustion, perhaps to better educate and mobilize participation from the research community.

A.3 Questions concerning the resulting portfolio of awards under review.

RESULTING PORTFOLIO OF AWARDS	APPROPRIATE, NOT APPROPRIATE[3], OR DATA NOT AVAILABLE
5. Overall quality of the research and/or education projects supported by the program. Comments: Appropriate. The quality of projects funded by the CT program is very high. The majority of the funded proposals received ratings of Very Good or Excellent from their reviewers, even though the CNS research community is quite critical.	
6. Does the program portfolio promote the integration of research and education? Comments: Appropriate. Most funded proposals tend to integrate education with research in a fairly "classic" manner, such as through the training of graduate students or the integration of research and results into both undergraduate and graduate courses. The COV believed this was a perfectly reasonable approach for many proposals to take. As well, we were pleased to find several examples of funded projects that had substantial or non-conventional educational outreach components.	

3. Are awards appropriate in size and duration for the scope of the projects?

Comments:

Data not available. We did not see specific data on average award size and duration within the Cyber Trust program specifically, but only CNS and CISE as a whole. However, we did note that average award size and duration for CNS was consistent with that of both CISE and the NSF as a whole.

4. Does the program portfolio have an appropriate balance of:

Innovative/potentially transformative projects?

Comments:

Appropriate. Based both on data in the self-study report, and on our sampling of funded proposals, we found a healthy and appropriate balance of innovative research ideas, proposals funded under programs explicitly designed to cultivate early or exploratory research (such as SGER), and high-risk proposals.

5. Does the program portfolio have an appropriate balance of:

Inter- and Multi- disciplinary projects?

Comments:

Appropriate.

6. Does the program portfolio have an appropriate balance considering, for example, award size, single and multiple investigator awards, or other characteristics as appropriate for the program?

Comments:

Appropriate. We were able to adequately form an opinion based on the self-study data supplemented with sampling of award jackets, but the COV recommends supplementing the self-study with additional data in the future. For example, distributional data on award sizes and durations, in addition to the averages provided, would give us a more detailed picture. Though not explicitly part of our charge, we were also curious about issues such as the degree to which average costs of personnel were increasing over time relative to average award sizes.

7. Does the program portfolio have an appropriate balance of:

Awards to new investigators?

NOTE: A new investigator is an investigator who has not been a PI on a previously funded NSF grant.

Comments:

Appropriate. The COV was surprised and pleased that a substantial fraction of awards within the CT program were granted to new PIs.

8. Does the program portfolio have an appropriate balance of:

Geographical distribution of Principal Investigators?

Comments:

Appropriate.

9. Does the program portfolio have an appropriate balance of:

Institutionnel types?

Comments:

Appropriate. (Note that to answer this question, the COV requested supplemental data not present in the self-study report, and not easily obtained by sampling jackets.)

10. Does the program portfolio have an appropriate balance:

Across disciplines and sub disciplines of the activity?

Comments:

Appropriate. There is no formal classification of projects into specific areas of CT, so our answer to this question is based on a qualitative assessment and general observations rather than statistics. Cyber Trust proposal solicitations are broad and inclusive, rather than narrowly focused, and as a result, the resulting portfolio of awards has a significant diversity of research topics, and an appropriate balance across disciplines within the Cyber Trust and security communities. Broad solicitations enable the community to pursue high-risk research, and define their own interests.

11. Does the program portfolio have appropriate participation of underrepresented groups?

Comments:

Appropriate. Participation of underrepresented groups is consistent across the clusters within CNS, and seemed consistent with the levels in the general CS community (for example, by comparing with CRA Taulbee survey data).

12. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.

Comments:

Appropriate. The COV commends the Cyber Trust program on the relevance of its funded portfolio. For example, the portfolio contains projects that cover each of the ten priority areas identified within the 2005 PITAC report on Cyber Security, entitled "Cyber Security: a Crisis of Prioritization", as well as the 22 research priority areas listed in the 2006 Federal Plan for Cyber Security and Information Assurance Research and Development, developed by the NSTC's Interagency Working Group on Cyber Security and Information Assurance.

13. Additional comments on the quality of the projects or the balance of the portfolio:

The Cyber Trust portfolio consists of high quality projects, and has an appropriate and vibrant diversity of topics, scales, risk/reward tradeoffs, and PIs.

A.4 Management of the program under review.

1. Management of the program. Comments: The Cyber Trust program is being managed effectively and competently, even though the program and its officers face significant resource constraints and the inevitable challenges of assembling panels with high quality reviewers. The COV spent a fair amount of time interacting with Karl Levitt, Program Officer for Cyber Trust, and was very impressed with his judgment, his ability to manage the internal and external pressures his program faces, and the degree to which he understands the research interests and needs of the Cyber Trust community of PIs.
2. Responsiveness of the program to emerging research and education opportunities. Comments: The Cyber Trust program solicitations have consistently contained a healthy balance between generality and specific areas of focus interest. The generality in the solicitations enable the community of PIs to pursue research of local interest, to find and engage in high-risk activities, and to be creative and pursue research in unanticipated, but important areas. Some of the language in the solicitations highlights areas of particular interest in detail; this plays the important role of focusing some of the community on problems of high national priority. The resulting balance is effective: areas of importance are highlighted and pursuing research on

them is encouraged and rewarded, but the community is not unduly hindered or channeled into narrow areas of interest. The COV commends the Cyber Trust program on this delicate but important balance, and encourages the program to invest deliberate effort to maintain and preserve it in the future.

Similarly, we found that the Cyber Trust program is receptive and responsive to educational opportunities within the computing security community, and that the program officer is aware of and encourages PIs to take full advantage of emerging opportunities when they present themselves.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

Comments:

The Cyber Trust portfolio is primarily reflective of the research interests of its constituent community, which we believe is appropriate. The program officer plays an active role within the community, has been successful at participating in and stimulating discussions about research priorities and opportunities at conferences and workshops, and has succeeded in reflecting both national and community interests in the program solicitations. One of the major tools the program officer can use to shape the portfolio is assembling qualified panels that are reflective of the interests and needs of the community, and ensuring that proposals are guided to appropriate panels; the officer has been using this tool appropriately and effectively. As a result, the program's portfolio, and the research, educational, and infrastructure priorities it reflects, is driven primarily by the community, with guidance -- but without undue constraints -- from the top.

There is the potential that some short-term relief of the immense research budgetary pressures faced by CISE will occur soon, and that Cyber Trust might benefit from this. If this does occur, the COV encourages the Cyber Trust director to be strategic in how it takes advantage of new funds, and to include the community in the planning process.

4. Responsiveness of program to previous COV comments and recommendations.

Comments:

The previous CNS COV did not specifically address the Cyber Trust program, so this question is not applicable.

5. Additional comments on program management:

PART B. RESULTS OF NSF INVESTMENTS

B. Please provide comments on the activity as it relates to NSF's Strategic Outcome Goals. Provide examples of outcomes ("highlights") as appropriate. Examples should reference the NSF award number, the Principal Investigator(s) names, and their institutions.

B.1 <u>OUTCOME GOAL for Discovery:</u> "*Foster research that will advance the frontier of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering."*

Comments:

The Cyber Trust portfolio delivered significant research and intellectual discoveries during FYs 2006-2008. We sampled the project highlights offered by the NSF itself, which contained many interesting and intellectually strong results, and also sampled project jackets and associated web sites directly, focusing in particular on the 10 largest awards from the last 3 years. The COV confirmed that the examined projects have been effective at conducting high-quality research and making interesting discoveries.

Of particular note are the following:

- *Award# 0433540, Michael Reiter, Carnegie Mellon University.* The "Center for Security Through Interaction Modeling" has produced many top tier publications in security and privacy, including publications at the IEEE Symposium on Information Security and Privacy, and at the ISOC Symposium on Network and Distributed Systems Security. Several discoveries reported in 2006 and beyond include automatic generation of worm signatures for static or self-modifying computer worms, and execution-based filtering of suspected malware. These results have been highly cited in the academic community, and have formed (and informed) the basis for many additional advances towards defense against malware.

- *Award# 0627513, Norman M. Sadeh, Carnegie Mellon University.* The "User-Controllable Security and Privacy for Pervasive Computing" project is developing new interfaces that combine user-centered design principles with dialogue, explanation and learning technologies to assist users in specifying and refining policies. It has conducted user studies on several applications, including instant messaging, directory service, and physical resource access control, on which policies have been implemented and tested from large user bases. This project is making discoveries towards an important and difficult problem, namely making the management of security usable and effective.

Award# 0716252, Douglas Blough, Georgia Institute of Technology. The "MedVault" project has produced many interesting results on a new approach to medical information protection, namely cross-layer identity and access management (IAM), in which access control and other security mechanisms are coordinated among the different information access layers in health care environments. Their approach leverages recent work on attribute-based and credential-based systems. The project has developed techniques for selective disclosure of signed information to allow patients to store information from a health care organization in their records, and then to selectively disclose this information to other health care organizations in a verifiable and tamper-proof manner. (The project research artifacts web site http://medvault.gtisc.gatech.edu/index.php?page=publications contains many interesting publications, presentations, as well as two demo videos.)

B.2 OUTCOME GOAL for Learning: *"Cultivate a world-class, broadly inclusive science and engineering workforce, and expand the scientific literacy of all citizens."*

Comments:

During FYs 2006-2008, several efforts in the Cyber Trust portfolio have made significant educational contributions, for the students at the nation's institutions, engineers and scientists in the workforce, and for an informed citizenry as well. The majority of awards we sampled

made “classic” and appropriate educational contributions by developing and teaching new courses, training and supervising student researchers, revising curriculum, and organizing and running workshops. Several projects in the portfolio went significantly above and beyond, containing outstanding and impactful educational contributions.

We outline two efforts in particular:

- *The Center for Correct, Usable, Reliable, Auditable, and Transparent Elections (Award# 0524252, Aviel Rubin, Johns Hopkins University).* The center has successfully informed and educated researchers, citizens and government officials alike on the dangers of faulty voting hardware and the fundamental challenges inherent in electronic voting systems. Student observers of the November 2006 election in Santa Clara and San Mateo counties in California wrote up a report documenting election procedures and their impact on correctness and verifiability that was submitted to the Election Assistance Commission. Several members of the project helped educate and inform designers of next-generation voting hardware. By testifying before election committees, members of the project have also educated the design of next-generation voting processes within local, state and federal governments. Members have also testified before election boards abroad, including during the presidential elections of Kazakhstan.

  Workshops organized by the center on the topic have expanded the outreach activities beyond the co-PIs and their organizations; these include a new recurring, peer-reviewed workshop on Electronic Voting Technologies. Finally, center activities have enjoyed continued presence in popular press, including the New York Times front page, which contributes greatly in educating and sensitizing the citizenry.

- *The TCIP project (Award# 0524695, William Sanders, University of Illinois at Urbana-Champaign).* produced three interactive Java-based educational activities on (1) power and energy in the home, (2) the power grid, and (3) power economics and emissions. They are especially geared towards teaching middle-school and high-school students about the critical topics of power, energy, and the national power grid. Additionally, 30 graduate students and 7 undergraduate students worked on TCIP topics in 2007-2008. The project web site also shows summer schools and internships that were held in conjunction with the project; see: http://www.iti.illinois.edu/content/tcip-education .

B.3 <u>OUTCOME GOAL for Research Infrastructure: "</u>*Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools."*

Comments:

The Cyber Trust program has funded several quality efforts to provide infrastructure and tools for cyber security research. Two examples of invaluable testbeds and tools for security experimentation are:

- *The DETER testbed (Award number: 0335298, "EIN: Collaborative Research: Cyber Defense Technology Experimental Research Network," PI: B. Clifford Neuman, Information Sciences Institute, University of Southern California)*. DETER is a national-scale cyber security research testbed. The testbed provides a unique environment where government, academic, and industry cyber security researchers safely analyze attacks and develop attack mitigation and confinement strategies. The testbed (as well as the associated EMIST project) also provides useful tools and resources to enable comparable and repeatable benchmarking experiments, allowing researchers to build a coherent scientific base for cyber security experimentation. With its strong security, containment, and usage policies, the testbed has fulfilled a role that is not met by the other large-scale testbed facilities such as PlanetLab or Emulab. Information about the testbed can be accessed at <u>http://www.deterlab.net/</u>.

- *The TCIP project (Award number: 0524695, "CT-CS: Trustworthy Cyber Infrastructure for the Power Grid (TCIP)," PI: William Sanders, University of Illinois at Urbana-Champaign).* TCIP has constructed a testbed that includes power system hardware and software, as well as advanced simulation capabilities. The testbed is becoming a unique national resource for experimenting with next-generation infrastructure. The project web site (<u>http://www.iti.illinois.edu/content/tcip-trustworthy-cyber-infrastructure-power-grid</u>) gives information on the tools developed.

From our survey of new awards during the 2006-2008 period, focusing in particular among larger scale awards, the COV did note an apparent reduction in such infrastructure efforts. The committee asserts that the continued support of infrastructure projects is critical to supporting the needs of the research community.

PART C. OTHER TOPICS

C.1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

None

C.2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

none

C.3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

none

C.4. Please provide comments on any other issues the COV feels are relevant.

none

C.5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

The COV review process was well managed, but we identified a few areas of potential improvement:

- The time spend one-on-one with Karl Levitt, the CT Program Officer, was extremely valuable. We would recommend giving the COV earlier access to the program officers.

- It was somewhat difficult at the onset of the process for the COV to determine the scope of their duties, and to decide upon a reasonable plan for performing them. As a result, a lot of the first day was spent acclimatizing rather than being fully productive. It might be worth identifying cluster chairs earlier, and bringing them out before the rest of the COV so that the chairs can formulate a plan and the rest of the COV can hit the ground running.

- The self-study document was very useful, and it helped the COV answer many of the more quantitative or mechanical questions in the review template. There were additional questions in the review template that could have been similarly answered with pre-processed data, but the self-study document didn't contain the appropriate analysis (such as histograms of funding rates and durations per program, per year, funding success rates for new PIs vs. established PIs, success rates for underrepresented minorities vs. other, and perhaps a complete set of jackets for a full proposal solitication / panel / funding decision cycle).

Cluster IV Computing Research Infrastructure (CRI)

FY 2009 REPORT TEMPLATE FOR

NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV: May 11-13, 2009
Program/Cluster/Section: Cluster IV – Computing Research Infrastructure
Division: CNS
Directorate: CISE
Number of actions reviewed: 36 **Awards:** 20 **Declinations:** 16 **Other:** Assignment of proposals was from CRI program. Subcommittee requested additional actions from GENI initiative and MRI program.
Total number of actions within Program/Cluster/Division during period under review: Total number of actions were provided for entire CNS cluster, and were not available for CRI cluster (see Table 2 of CNS COV self-study report) **Awards:** **Declinations:** **Other:** Total numbers of actions from GENI initiative and from MRI program were not available.
Manner in which reviewed actions were selected: - Subcommittee made a random selection from the provided samples of 73CRI actions. - For GENI, specific requests where made to PO. - For MRI, samples were requested from PO.

PART A. INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

A.1 Questions about the quality and effectiveness of the program's use of merit review process.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE[1]
1. Are the review methods (for example, panel, ad hoc, site visits) appropriate? Yes. The review process in this cluster is primarily conducted using a panel of expert reviewers. In a small number of cases, ad hoc mail reviews are also used. The balance seemed appropriate. A small number of proposals might have benefited from a site visit, but current travel funding rules make such visits difficult. We suggest that CNS reexamine travel funding rules related to site visits.	
2. Are both merit review criteria addressed In individual reviews? • Overall standard of individual reviews was good, with some exceptions, in particular some panelists provided shallow reviews that added little to the review process. • The intellectual merit was adequately addressed in most reviews. • In approximately 15% of reviews, broader impact was addressed inadequately, with superficial coverage. For example, stating that a project "will enhance student education" or will involve a particular underrepresented group was found to be too generic. • In some of the jackets, it appeared that panelists had been given more concrete guidance on how to address the review criteria, and we applaud this idea. In panel summaries? • Intellectual merit was adequately addressed in most cases. • About 10% of panel summaries failed to adequately address the broader impact criteria. • The COV suggests the following strategies to enhance quality of reviews: o Educate panelists/reviewers on desired quality and specificity. o Examine strategies that can limit impact of bad "outlier" reviews on funding decisions. o Have reviewers state their level of expertise with topics that they review	

In Program Officer review analyses?

- Overall, the PO analyses were very impressive and contained detailed well-reasoned explanations for recommendation based on the merit review criteria.

In one of the reviews we examined, there was evidence that reviewers had been given a checklist of issues to consider in examining intellectual merit and broader impacts. Although in the examined case, the checklist had been answered rather superficially, providing this kind of guidance to reviewers can help guide their thinking during the review process and should be encouraged. We have reproduced the checklist below to assist CNS and CISE in developing a broadly applicable checklist.

Intellectual merit checklist:

- How important is the proposed activity to advancing knowledge and understanding within its own field or across different fields?
- How well qualified is the proposer to conduct the project?
- To what extent does the proposed activity suggest and explore creative and original concepts?
- How well conceived and organized is the proposed activity?
- Are there sufficient resources?
- Is the proposal too expensive to be funded by individual grants?
- Will the infrastructure enable research that would not otherwise be possible?
- Is the infrastructure adequately described?
- Are the research and educational activities that it will support adequately described?
- Is the case made that the grant will enable new achievements that could not be done otherwise?

Broader impact checklist:

- How well does the activity advance discovery and understanding while promoting teaching, training and learning?
- How well does the proposed activity broaden the participation of underrepresented groups?
- To what extent will it enhance the infrastructure for research and education?
- Will the results be disseminated broadly to enhance scientific and technological understanding?
- What may be the benefits of the proposed activity to society?
- Are synergies among activities and participants demonstrated?
- Does the activity produce leverage, enable new sources of support, increase research recognition, industry participation, or links between institutions?
- Does the activity expand the number of individuals that can contribute to the research base?
- Are the management plans sound?
- Is the proposed infrastructure to be shared among research groups and/or useful to a broad research and educational community?

[1] If “Not Applicable” please explain why in the “Comments” section.

3. Do the individual reviewers provide substantive comments to explain their assessment of the proposals? About 90% of reviewers provide substantive comments, which clearly justify their review ratings.

4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)?

Yes. Overall, for both awarded and declined proposals, the panel summaries provided an informative explanation of the panel consensus and cover the rationale of the decision. In some cases, most noticeably for declined proposals, the summary didn't entirely capture the rationale in a way that would be helpful to PIs.

5. Does the documentation in the jacket provide the rationale for the award/decline decision?

(Note: Documentation in jacket usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), program officer review analysis, and staff diary notes.)

Yes. The review analyses provided a clear rationale for award/decline decisions, as well as the rationale for funding reductions.

6. Does the documentation to PI provide the rationale for the award/decline decision?

(Note: Documentation to PI usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), and, if not otherwise provided in the panel summary, an explanation from the program officer (written or telephoned with diary note in jacket) of the basis for a declination.)

Yes.

7. Is the time to decision appropriate?

Note: Time to Decision --NSF Annual Performance Goal: **For 70 percent of proposals, inform applicants about funding decisions within six months of proposal receipt or deadline or target date, whichever is later**. The date of Division Director concurrence is used in determining the time to decision. Once the Division Director concurs, applicants may be informed that their proposals have been declined or recommended for funding. The NSF-wide goal of 70 percent recognizes that the time to decision is appropriately greater than six months for some programs or some individual proposals.

Yes. Over the 3-year CoV period, approximately 85% of submitted proposals received a decision within a six-month period. This commendable performance exceeds NSF performance goals.

8. Additional comments on the quality and effectiveness of the program's use of merit review process:

There is a need to educate the reviewer community on reviewing best practices, such as timeliness of reviews, specificity of comments, value of broader impacts, etc..

A.2 Questions concerning the selection of reviewers.

SELECTION OF REVIEWERS	YES , NO, DATA NOT AVAILABLE, or NOT APPLICABLE[2]
1. Did the program make use of reviewers having appropriate expertise and/or qualifications? As far as the subcommittee could tell from the available data, yes. We recommend that the POs examine and document the level of expertise of reviewers in the topics reviewed to make future evaluation of this question more tractable.	
2. Did the program use reviewers balanced with respect to characteristics such as geography, type of institution, and underrepresented groups? Yes, POs have done a commendable job in this area. Note: Demographic data is self-reported, with only about 25% of reviewers reporting this information.	
3. Did the program recognize and resolve conflicts of interest when appropriate? Yes. Prospective reviewers are given access to full proposals prior to attending the NSF panel and are asked to verify whether there is a potential COI. The COI policy is again reviewed thoroughly before the panel begins proposal reviews. Panelists in conflict with a particular proposal do not participate in any aspect of reviewing the proposal.	

A.3 Questions concerning the resulting portfolio of awards under review.

RESULTING PORTFOLIO OF AWARDS	APPROPRIATE, NOT APPROPRIATE[3], OR DATA NOT AVAILABLE
7. Overall quality of the research and/or education projects supported by the program. Based on the expert reviews and awarded projects, the overall quality of the research and education projects are rated as very-good to excellent.	
8. Does the program portfolio promote the integration of research and education? Yes. This cluster provides a good mix of infrastructure development and acquisition at small, medium and large scale to support integration of research and education.	

3. Are awards appropriate in size and duration for the scope of the projects?

It was clear from reviewing jackets that POs have done a commendable job in a difficult budgetary situation. In many instances, POs made cuts very reluctantly, aware of the impact that those cuts would have, but trying to balance cutting budgets with funding more proposals. CNS must continue to argue aggressively for funds to meet the needs of its research community.

4. Does the program portfolio have an appropriate balance of:

Innovative/potentially transformative projects?

Based on the sample, the portfolio comprises a large number of innovative projects. However, truly transformative/high risk projects were rare. The lack of transformative projects may be due to all involved (researchers/scientists/NSF/panelists) attempting to guarantee a successful solution. This leads to awarding conservative, incremental (as opposed high risk) projects.

5. Does the program portfolio have an appropriate balance of:

Inter- and Multi- disciplinary projects?

Yes. The sample proposals each had some form of interdisciplinary and/or multidisciplinary components. In some cases, this involved sub-disciplines, e.g., networking and HPC, within an area. Other cases involved multiple disciplines, e.g., computer science and oceanography.

6. Does the program portfolio have an appropriate balance considering, for example, award size, single and multiple investigator awards, or other characteristics as appropriate for the program?

Research infrastructure projects typically involve multiple investigators.

7. Does the program portfolio have an appropriate balance of:

Awards to new investigators?

NOTE: A new investigator is an investigator who has not been a PI on a previously funded NSF grant.

Based on the self-study report (see Table 14), we concluded that approximately 7.5% of awards over the 3-year period involved new PIs. The self-study report did not list the number of unsuccessful applicants that would have been new investigators, making it difficult to fully answer this question.

8. Does the program portfolio have an appropriate balance of:

Geographical distribution of Principal Investigators?

Yes. POs have done a commendable job to achieve a balance in this category.

9. Does the program portfolio have an appropriate balance of:

Institution types?

From a random sampling of jackets, the balance seemed appropriate. The data in the self-study report only showed the balance with regard to MSIs. Here, the balance seemed roughly appropriate, e.g., in 2006, 9.2% of funded proposals in CRI were from MSIs, whereas in the total pool of CNS proposals, 10.7% of funded proposals were from MSIs.

10. Does the program portfolio have an appropriate balance:

Across disciplines and sub disciplines of the activity?

Yes.

11. Does the program portfolio have appropriate participation of underrepresented groups?

Yes. For example, from the 2006 data of in the self-study report, we found that about 20% of the applicant pool self-identified as women, and 20% of funded proposals had women PIs. Similarly, about 7% of applicants were minorities and about 7.9% of funded proposals were from minority PIs.

12. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.

Yes. The program is clearly serving the infrastructure needs regarding research and education, with a long-term vision. The GENI initiative has received national and international attention within the broad computer science community, and will make a significant contribution, both directly and indirectly, in addressing national priorities for the future generation of networks.

The MRI program is especially notable for having a good track record of generating high profile "nuggets" for CNS and CISE.

13. Additional comments on the quality of the projects or the balance of the portfolio:

The RI programs are serving a critical need of the research community in a commendable way.

There may be opportunities for the size of the MRI portfolio to increase within CNS. The CNS community needs to become more aware of the MRI program, and needs to learn how to better articulate its research instrumentation needs to a Foundation-wide audience.

A.4 Management of the program under review.

1. Management of the program.

On the whole, POs seem to be doing an excellent job. The subcommittee did not have access to sufficient information to fully assess the management practices of CNS. As a general point, we note that management experience gleaned in academia may not always match the management roles within CNS and CISE. The self study did not explain how management best practices are communicated as POs rotate in and out, nor how CNS or CISE evaluates the effectiveness of its management strategies. We can, however, recommend that CNS (and CISE as whole) determine and adopt appropriate management best practices at all levels, and apply sufficient scrutiny to ensure that these best practices work well. We also recommend that information about management practices and their evaluation be included in future self studies.

The subcommittee emphasizes that empowerment of Program Officers is a great asset of NSF that must be continuously fostered and guarded.

GENI has a unique management structure that may not be well understood by the wider networking research community. The management processes and practices in CNS examined by the COV were found to be proper and thorough. Greater transparency is required with regard to the review process of subcontracts managed by BBN.

2. Responsiveness of the program to emerging research and education opportunities.

Excellent. This cluster is very responsive to emerging research and educational opportunities. In fact, the POs do an excellent job of encouraging PIs to pursue various programs and also offer substantive advice on how best to position a proposed project within the context of the overall science and engineering community.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

The planning and prioritization process is open and welcomes the entire CISE community. The process is conducted primarily via workshops, advisory meetings, and study groups.

The subcommittee did not see CNS-internal documents articulating priorities for fostering new research directions.

4. Responsiveness of program to previous COV comments and recommendations. We recommend that a specific action plan be formulated to respond to COV comments and recommendations with evaluation mechanisms to determine how well the plan has been followed.
5. Additional comments on program management:

PART B. RESULTS OF NSF INVESTMENTS

B. Please provide comments on the activity as it relates to NSF's Strategic Outcome Goals. Provide examples of outcomes ("highlights") as appropriate. Examples should reference the NSF award number, the Principal Investigator(s) names, and their institutions.

B.1 OUTCOME GOAL for Discovery: "*Foster research that will advance the frontier of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering."*

A significant portion of the funded research infrastructure project portfolio is at the cutting edge of science and engineering.

For example, in the award 0454233/0627822, researchers at UC San Diego and SUNY at Buffalo develop a toolbox for automatic recognition of facial expressions. A discovery of the project is that computers can recognize facial expression better than humans can. The developed system is able to differentiate real from faked pain expressions, and has detected driver drowsiness with very high accuracy. Projects 0224363 and 0420836 at the University of Minnesota and Berea College have developed an infrastructure for swarms of miniature mobile robots. A number of these robots were reportedly deployed in reconnaissance and recovery missions in Afghanistan and Iraq.

B.2 OUTCOME GOAL for Learning: "*Cultivate a world-class, broadly inclusive science and engineering workforce, and expand the scientific literacy of all citizens."*

The overwhelming majority of funded infrastructure projects that were looked at by the subcommittee were found to contain components for training future scientists and engineers.

B.3 OUTCOME GOAL for Research Infrastructure: "*Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools."*

The portfolio of research infrastructure projects collectively builds the nation's research

infrastructure in computer science research and education. Instrumentations of infrastructures consist not only of acquisitions of advanced technology. In addition, many unique infrastructures are being built within the context of funded research infrastructure projects.

As an example, the project 0544123 has developed a one of a kind test facility at Pennsylvania State University to study the impact of radiation-induced errors on computer chips. Understanding the impact of such errors, whose frequency is anticipated to increase as the size of chips decreases, and developing methods for error analysis and mitigation is of critical importance in safety critical systems such as aircrafts and nuclear reactors.

The instrumentation and infrastructure of the GENI initiative may have a transformational impact by addressing a nationally critical need for next generation communication networks that are significantly advanced compared to the currently deployed Internet infrastructure.

PART C. OTHER TOPICS

C.1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

The RI cluster has a well managed, rich portfolio of projects. The proportion of high-risk and innovative projects is appropriate. The subcommittee did not find deficiencies or gaps that need to be highlighted at this point.

The establishment of research infrastructures is critical as they are a critical enabling factor for conducting leading research that addresses national priorities and needs in computer and network systems.

The programs of this cluster are well-integrated across CISE and across the Foundation in terms of cross-cutting activities and joint funding. This should be strengthened and nurtured, for example, through additional interagency funding.

C.2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

The program is meeting and exceeding all program specific goals. The committee commends the POs on the efficiency and timeliness of the review process. The program is highly successful in including all parts of the broader community, ranging from small colleges, minority serving institutions, to PhD granting research institutions.

C.3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

The CNS community of researchers could greatly benefit from additional guidance on preparing successful proposals for agency-wide cross-cutting programs.

C.4. Please provide comments on any other issues the COV feels are relevant.

Programs such as GENI have such a large scope and potential of transformational impact that broader involvement of the community in and beyond computer science should be fostered.

C.5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

We suggest to improve the efficiency of the COV meeting by a better organization and structure put in place in advance of the meeting, and by providing materials and guidance to COV members prior to the meeting. Reading the self-study report and previous COV report prior to the meeting will avoid stepping through a learning-curve during the meeting. The high-level presentations by POs on Day 1 were less valuable than one-on-one meetings with POs in the cluster subcommittees. Our subcommittee found the meetings with POs very helpful and valuable.

We recommend that CNS and CISE record these best practices, and others learned from prior COVs, and take steps to ensure that all COVs are able to work productively from day one.

In addition, we recommend greater continuity between successive CoV visits. If a suitably chosen subset of the previous CoV is invited back for the next one, best practices are more likely to be developed and retained, and any unaddressed issues from the previous visit are more likely to be revisited.

Finally, we note that the self study mostly concerned itself with questions in the “A” group, but we see no reason at all why it could not also have examined questions from the “B” and “C” groups.

Cluster V Computer Education & Workforce (EWF)

FY 2009 REPORT TEMPLATE FOR

NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV:
Program/Cluster/Section:
Division:
Directorate:
Number of actions reviewed: **Awards:** **Declinations:** **Other:**
Total number of actions within Program/Cluster/Division during period under review: **Awards: 44** **Declinations: 40** **Other:**
Manner in which reviewed actions were selected:

PART A. INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

A.1 Questions about the quality and effectiveness of the program's use of merit review process.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE[1]
1. Are the review methods (for example, panel, ad hoc, site visits) appropriate? Yes Comments: In general, the process used seems to be appropriate. However, for the proposals relevant to the cluster, of the two in our sample where we expected to see a site visit report, only one had evidence of a site visit. Otherwise, the material presented to us seems to be within NSF guidelines.	
2. Are both merit review criteria addressed In individual reviews? No. Missing in a significant percent. For example, the broader impact criteria was missing or poorly addressed in some cases. In panel summaries? The panel summaries provided a more complete picture than some of the reviews. If the proposal was not competitive this was often found to be well stated in these summaries. In Program Officer review analyses? Occasionally, the review criteria were not addressed. It was occasionally difficult to discern what the problem with a proposal was from these reports. In some cases they appeared to be formulaic. Comments:	

[1] If "Not Applicable" please explain why in the "Comments" section.

3. Do the individual reviewers provide substantive comments to explain their assessment of the proposals? A significant minority of reviews lacked substance. Comments:
4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)? Yes. Cluster members looked through these and agreed that the panel summaries were commendable. Comments:
5. Does the documentation in the jacket provide the rationale for the award/decline decision? (Note: Documentation in jacket usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), program officer review analysis, and staff diary notes.) Yes. Comments:

6. Does the documentation to PI provide the rationale for the award/decline decision?

(Note: Documentation to PI usually includes context statement, individual reviews, panel summary (if applicable), site visit reports (if applicable), and, if not otherwise provided in the panel summary, an explanation from the program officer (written or telephoned with diary note in jacket) of the basis for a declination.)

Yes, in general. Sometimes the information was brief.

Comments:

7. Is the time to decision appropriate? Yes.

Note: Time to Decision --NSF Annual Performance Goal: **For 70 percent of proposals, inform applicants about funding decisions within six months of proposal receipt or deadline or target date, whichever is later**. The date of Division Director concurrence is used in determining the time to decision. Once the Division Director concurs, applicants may be informed that their proposals have been declined or recommended for funding. The NSF-wide goal of 70 percent recognizes that the time to decision is appropriately greater than six months for some programs or some individual proposals.

Comments:

8. Additional comments on the quality and effectiveness of the program's use of merit review process:

Reviews could be more comprehensive and address both merit review criteria.

It was not clear whether a multidisciplinary review panel was used to evaluate proposals that had multidisciplinary components.

A.2 Questions concerning the selection of reviewers.

SELECTION OF REVIEWERS	YES , NO, DATA NOT AVAILABLE, or NOT APPLICABLE[2]
1. Did the program make use of reviewers having appropriate expertise and/or qualifications? Yes, but see comment on multi-disciplinary reviewers (A.1.8) Comments:	
2. Did the program use reviewers balanced with respect to characteristics such as geography, type of institution, and underrepresented groups? Yes. We looked carefully at three programs and saw a range of representation. Note: Demographic data is self reported, with only about 25% of reviewers reporting this information. Comments:	
3. Did the program recognize and resolve conflicts of interest when appropriate? The process that is in place appears adequate. Comments:	

Additional comments on reviewer selection:

An effort is being made to get a range of reviewers and we recognize that this is a difficult process. Individual Program Directors appear to make efforts to be representative.

A.3 Questions concerning the resulting portfolio of awards under review.

RESULTING PORTFOLIO OF AWARDS	APPROPRIATE, NOT APPROPRIATE[3], OR DATA NOT AVAILABLE
9. Overall quality of the research and/or education projects supported by the program. Very good quality education and research experiences for projects in this cluster and very promising for the future of computer science. Comments: The REU awards were diverse in terms of participants and topics. All topics that were reviewed were STEM, with primary science and cutting edge research. BPC projects appear to be good. There are novel interventions with intellectual merit. CPATH has made diverse awards that appear promising.	
10. Does the program portfolio promote the integration of research and education? Yes. Comments: REU provides opportunities for the integration of research and education. Other programs in the cluster also support projects that have REU-like activities.	

3. Are awards appropriate in size and duration for the scope of the projects?

See comments.

Comments:

Successful REU projects may not be renewed because there is not enough funding to both continue existing ones and begin new ones. Expansion of funding for renewals of successful programs and establishment of new programs seems appropriate.

BPC awards are appropriate in size and duration for the scope. Alliances have site reviews before receiving an additional two years of funding. Looking forward to the future, the BPC program appears poised to continue this accountability in the results required for alliance extensions. BPC demonstration grants vary in size and scope, but appear appropriately funded.

There appear to be a wide range of CPATH award sizes.

4. Does the program portfolio have an appropriate balance of:

Innovative/potentially transformative projects?

Yes.

Comments:

For REU's the topics of the projects are varied. Many are very innovative and some actually will provide the participants with an introduction into potentially transformative science.

The BPC program appears to have tremendous value for transformation. There is potential for

innovation in BPC programs to propagate out to other initiatives, such as CPATH. Although the level of innovation seems high, the program does not appear high-risk.

CPATH has made strides in developing innovative approaches and building community to allow the approaches to become transformative.

5. Does the program portfolio have an appropriate balance of:

Inter- and Multi- disciplinary projects?

Yes.

Comments:

BPC and CPATH have many multidisciplinary programs, for example, artistic performance and computing.

For REUs many were inter-disciplinary. We observed that since REU sites are reviewed by division, truly multidisciplinary REU site proposals might be at a disadvantage since they may not have a champion within more narrowly focused panels.

6. Does the program portfolio have an appropriate balance considering, for example, award size, single and multiple investigator awards, or other characteristics as appropriate for the program?

Yes

Comments:

For CPATH there seems to be diversity along several axes of measurement.

7. Does the program portfolio have an appropriate balance of:

Awards to new investigators?

Yes.

NOTE: A new investigator is an investigator who has not been a PI on a previously funded NSF grant.

Comments:

The CPATH, REU and BPC programs appear to have brought in a fair number of new proposal writers to the NSF community.

8. Does the program portfolio have an appropriate balance of:

Geographical distribution of Principal Investigators?

Yes.

Comments:

9. Does the program portfolio have an appropriate balance of:

Institutional types?

Yes.

Comments: Selection of REU sites at larger institutions is appropriate given the nature of the program.

10. Does the program portfolio have an appropriate balance:

Across disciplines and sub disciplines of the activity?

Yes.

11. Does the program portfolio have appropriate participation of underrepresented groups?

Comments:

These programs appear to have broad representation from women and other minority PIs, minority serving institutions, and EPSCoR states.

12. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.

Yes.

The activities of the educational and workforce cluster are relevant to the Learning objective of the National Science Foundation at the highest level. These programs support national priorities such as workforce development, the America Competitiveness initiative, improving STEM education.

Finally, these programs have the potential to contribute to the future of the discipline by preparing the next generation of computer scientists.

Comments:
13. Additional comments on the quality of the projects or the balance of the portfolio:

A.4 Management of the program under review.

1. Management of the program. The BPC program appears to be exceptionally well managed. It has a well-articulated mission and it builds community to advance program goals. The program director does not rely exclusively on award instruments to advance the mission, but advances the mission through a broad range of activities and even provides additional support through PI meetings that multiply and expand the impact of individual projects. Other programs across CISE might benefit from the community building model of BPC. Comments:
2. Responsiveness of the program to emerging research and education opportunities. These programs are creating new education opportunities. They are in tune with the needs of the community. The BPC program has appropriately adjusted its mix of proposal size and focus to reflect the needs of a maturing community. For example, just as the first round of demonstration grants were expiring, a new track, LSA, was introduced to help those projects (and others) to expand their impact.

Comments:

11. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

We did not have any information with respect to this item.

Comments:

4. Responsiveness of program to previous COV comments and recommendations.

Previous recommendations to this cluster were to increase the number of proposals from minority serving institutions and to increase the number of large-scale, high-risk projects funded. As a whole, this cluster has improved performance in both of these areas through BPC. Both BPC and CPATH have implemented proposal writing workshops for new PIs from MSIs that helped to meet these recommendations.

Comments:

5. Additional comments on program management:

None.

PART B. RESULTS OF NSF INVESTMENTS

B. Please provide comments on the activity as it relates to NSF's Strategic Outcome Goals. Provide examples of outcomes ("highlights") as appropriate. Examples should reference the NSF award number, the Principal Investigator(s) names, and their institutions.

B.1 OUTCOME GOAL for Discovery: "*Foster research that will advance the frontier of knowledge, emphasizing areas of greatest opportunity and potential benefit and establishing the nation as a global leader in fundamental and transformational science and engineering."*

Comments:

While the programs in this cluster are not directly research focused, they advance the frontier of knowledge in two important ways. First, they advance education and workforce development in the computing disciplines. Second, they prepare a new generation of researchers from among students and other members of the population who benefit from the programs in this cluster.

B.2 OUTCOME GOAL for Learning: "*Cultivate a world-class, broadly inclusive science and engineering workforce, and expand the scientific literacy of all citizens."*

Comments:

REU should have a huge impact on a world-class workforce.

Examples from the BPC include:

AccessComputing Alliance #0540615 which provides access to computing education for disabled students. (Richard Ladner, University of Washington).

Collaborative Research: BPC-DP: Improving Minority Student Participation in the Computing Career Pipeline with Culturally Situated Design Tools. (#0634329, Eglash, Rensselaer Polytech Institute)

CPATH example:

0722327 (Leonard Pitt, U of IL, Urbana-Champaign) Informatics and Computation Ubiquitous throughout Undergraduate Education indicates an expansion of informatics across several disciplines.

B.3 <u>OUTCOME GOAL for Research Infrastructure: "</u>*Build the nation's research capability through critical investments in advanced instrumentation, facilities, cyberinfrastructure and experimental tools."*

N/A

Comments:

PART C. OTHER TOPICS

C.1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

The CPATH program is newer and appears to have a great deal of potential. Community building in this context is essential and should be encouraged.

There should adequate opportunity for new and high-quality renewing REU site awards to co-exist.

C.2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

Cluster programs relate to education and workforce development and have broad impact. Cluster programs promote NSF's mission with respect to "Learning", developing the S&E workforce and scientific literacy among the broadest possible spectrum of society. These programs should receive the highest level of support within NSF.

C.3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

One agency wide issue is the enforcement of the two merit review criteria. Some reviewers prepare reviews that omit one of the criteria. A modified charge to the reviewers and an emphasis on the fact that both criteria are important would be helpful.

C.4. Please provide comments on any other issues the COV feels are relevant.

None.

C.5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

Materials were not readily available until right before the meeting. It would be helpful to have time to absorb some of the materials and thus arrive with a foundation for work.

Not everyone knew which cluster they would be assigned to and this made it difficult to prepare by becoming familiar with the programs in the cluster ahead of time.

The process could be structured better to make more efficient use of time:

1. Know which cluster you belong to ahead of time to give sufficient time to review the previous COV with particular attention to that cluster and relevant programs.
2. Notify the chair of the cluster well ahead of time and be sure they agree to chair the cluster. Inform and engage the chair ahead of time regarding the process and be sure that the chair can make the time commitment required.
3. At the initial gathering, review the template and indicate how the self-study relates to answering those questions in part A.
4. In clusters, provide time to review the pamphlets about each program—these were handed out at the beginning of the meeting with little or no time for review.
5. Provide structured interview time with each program director, individually.